

Fig. 1a

Fig. 1b

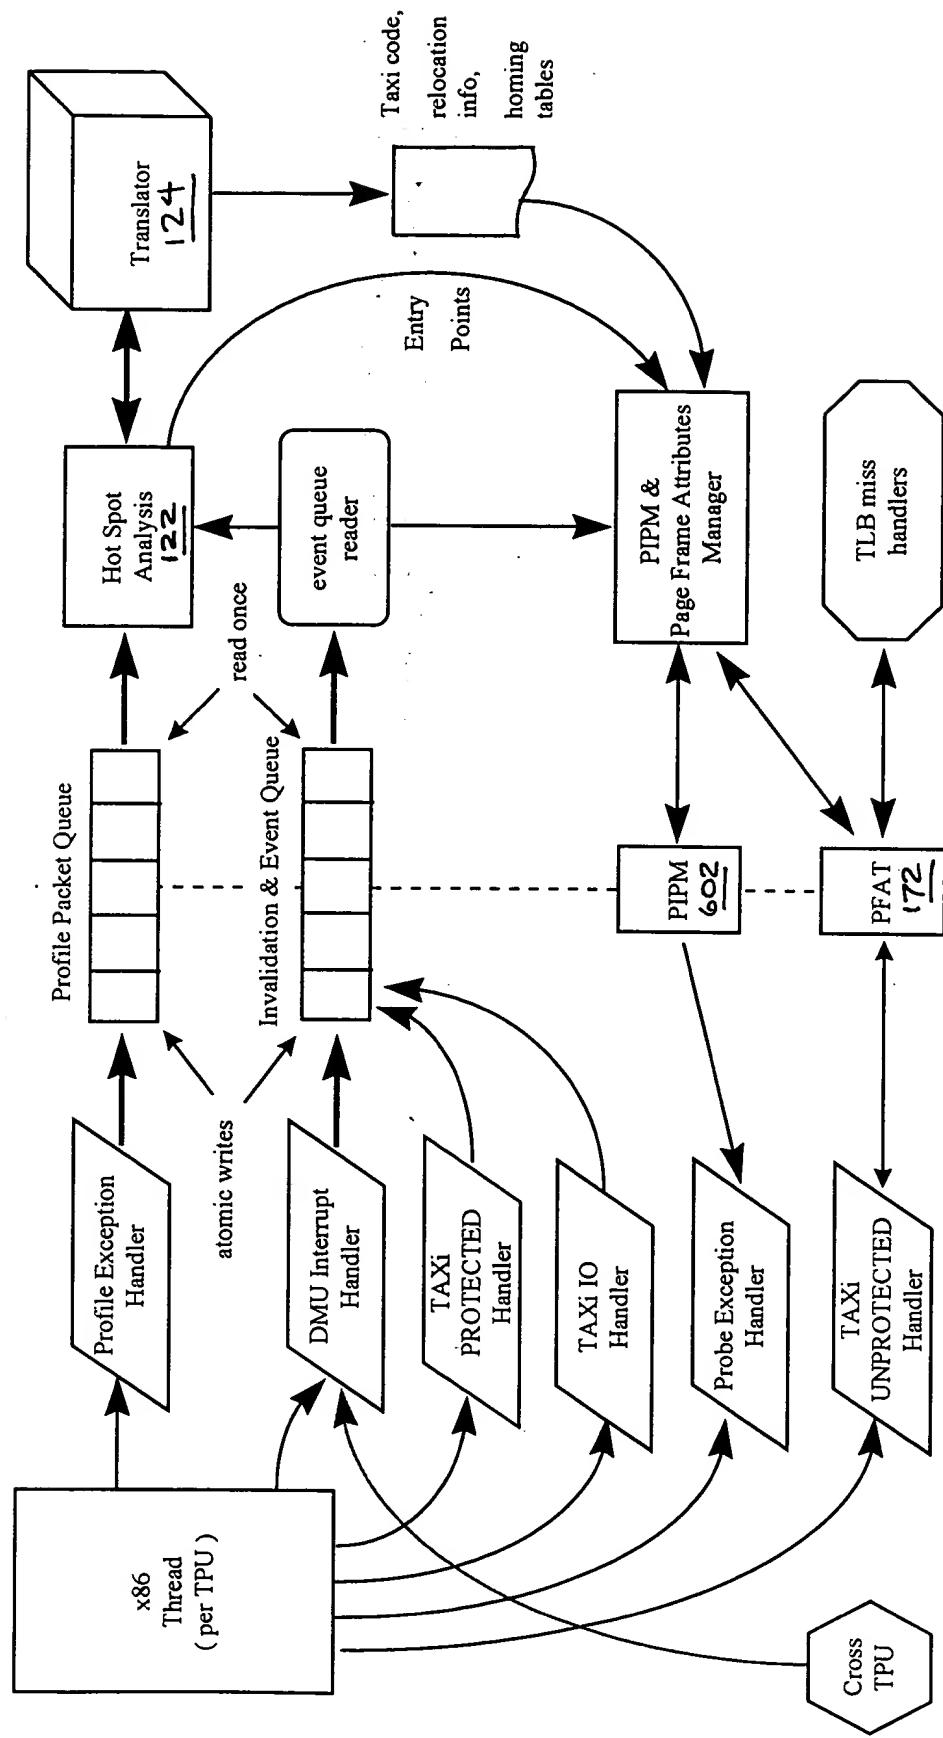
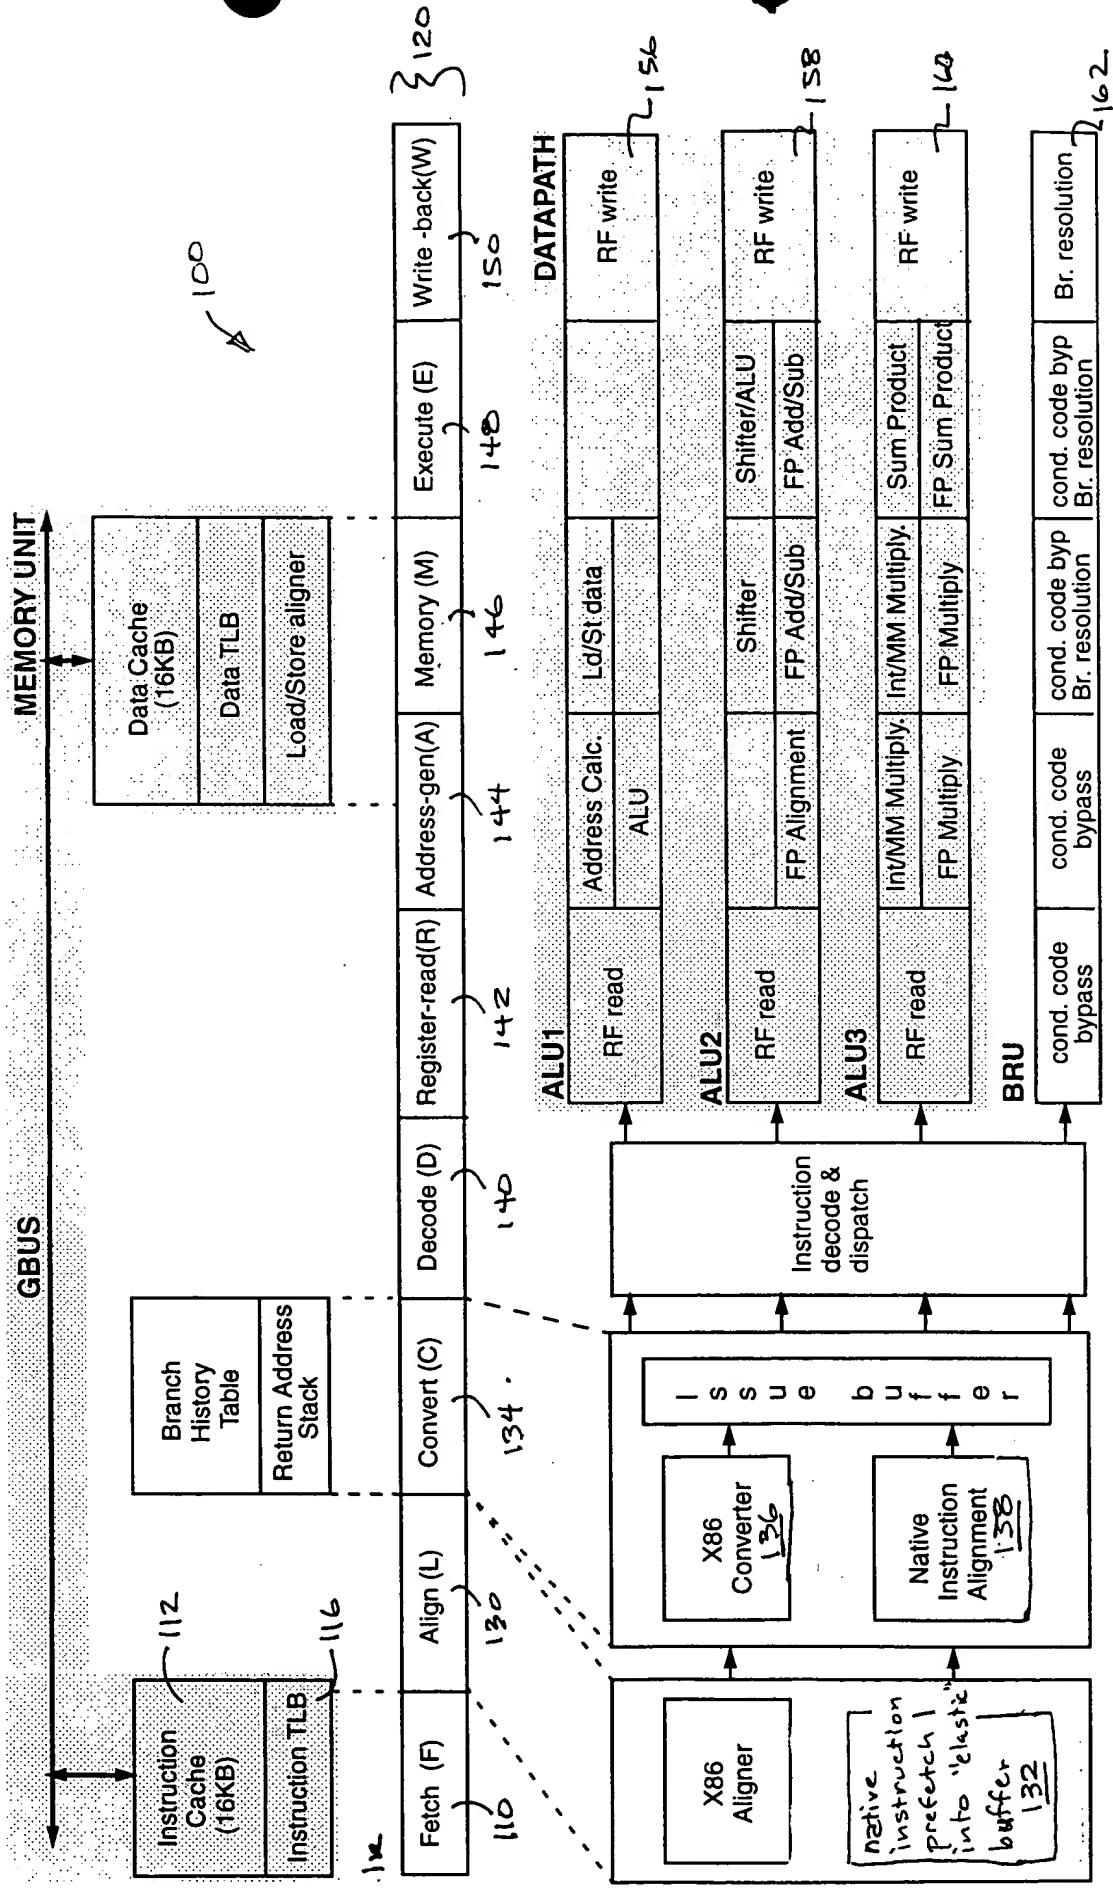


Fig. 1b

Fig. 1c



50 SHEETS  
22-141 22-142 22-144  
100 SHEETS  
200 SHEETS



Logical Address

Segment | offset

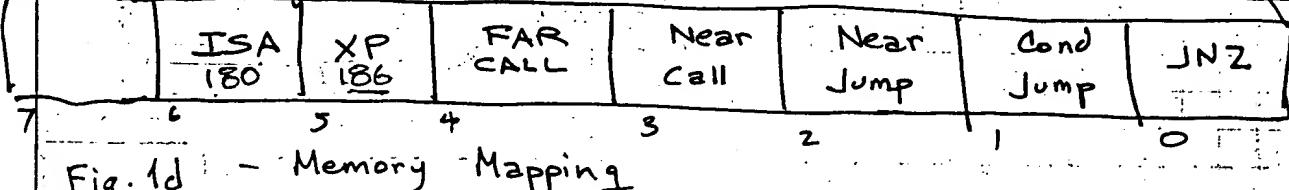
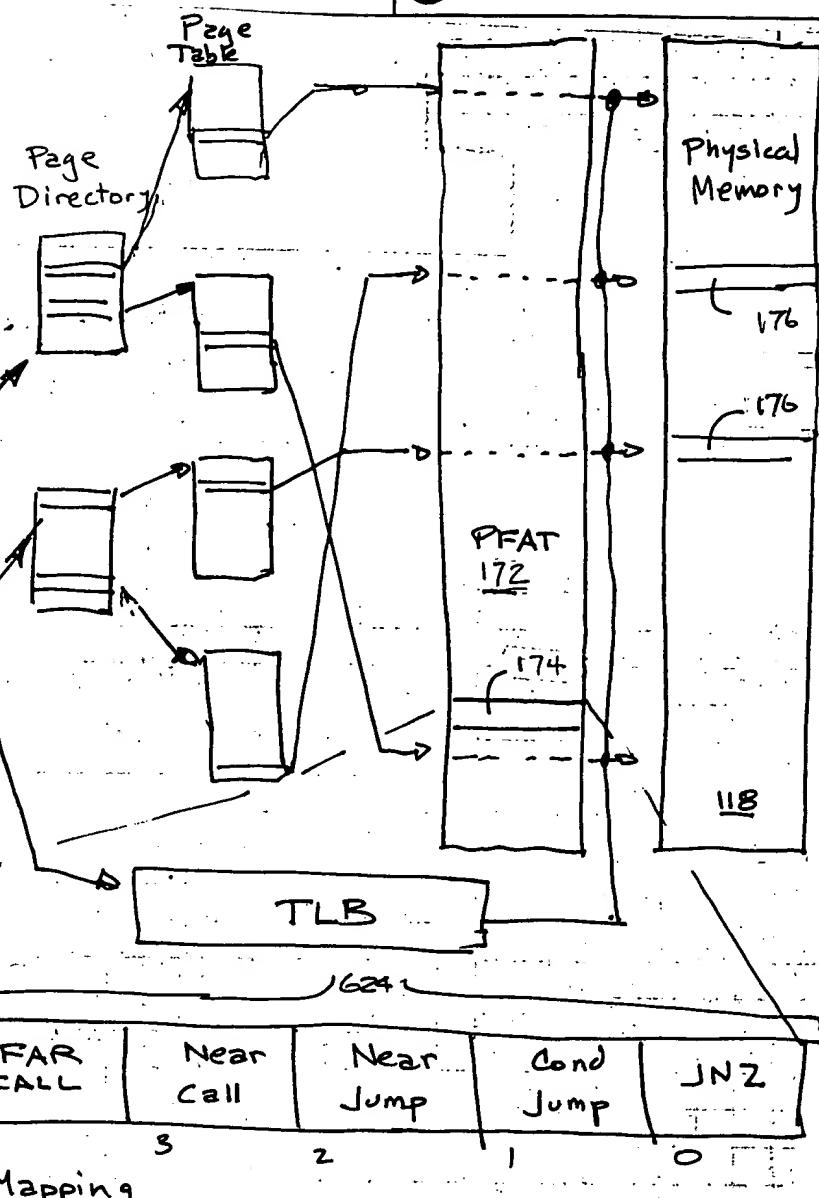
X86  
segment  
translation

Linear Address  
or  
Virtual Address

170 ↗

174 ↗

Fig. 1d : - Memory Mapping



22-141 50 SHEETS  
22-142 100 SHEETS  
22-144 200 SHEETS

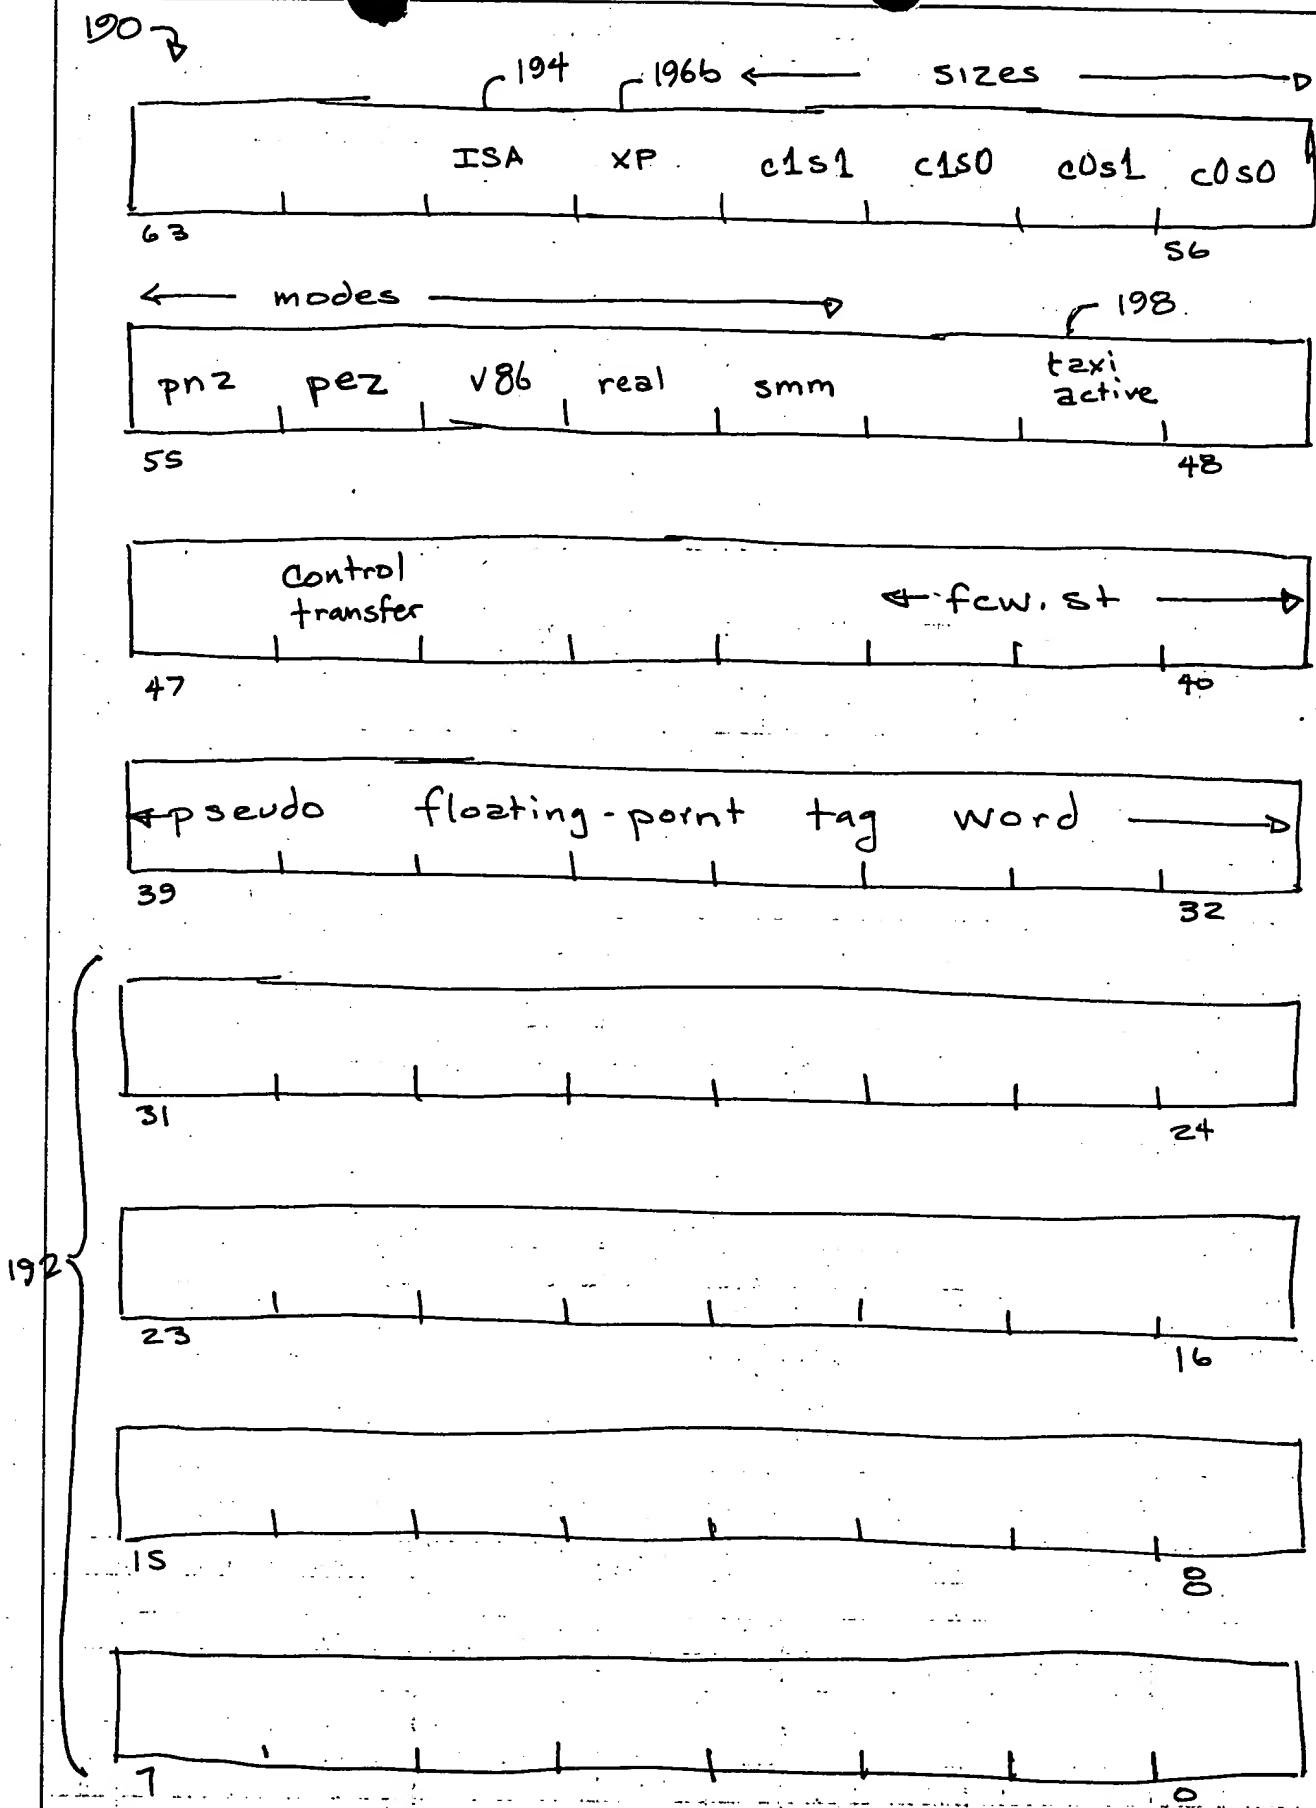


Fig. 1e

I-TLB property bits	Decoded property values			Interpretation	Instructions sent to:	Collect profile trace-packets?	Probe for translated code	I/O memory reference exceptions
	ISA	CC	Protected					
00	Tap	Tap	no	Native code observing native RISCy calling conventions	Native decoder	No	No	Fault if SEG.tio
01	Tap	x86	no	Native code observing x86 calling conventions	Native decoder	No	No	Fault if SEG.tio
10	x86	x86	no	x86 code, unprotected - TAX! profile collection only	x86 HW converter	If enabled	No	Trap if profiling
11	x86	x86	yes	x86 code, protected - TAX! code may be available	x86 HW converter	If enabled	Based on I-TLB probe attributes	Trap if profiling

Fig. 2a : Significance of the I-TLB property bits

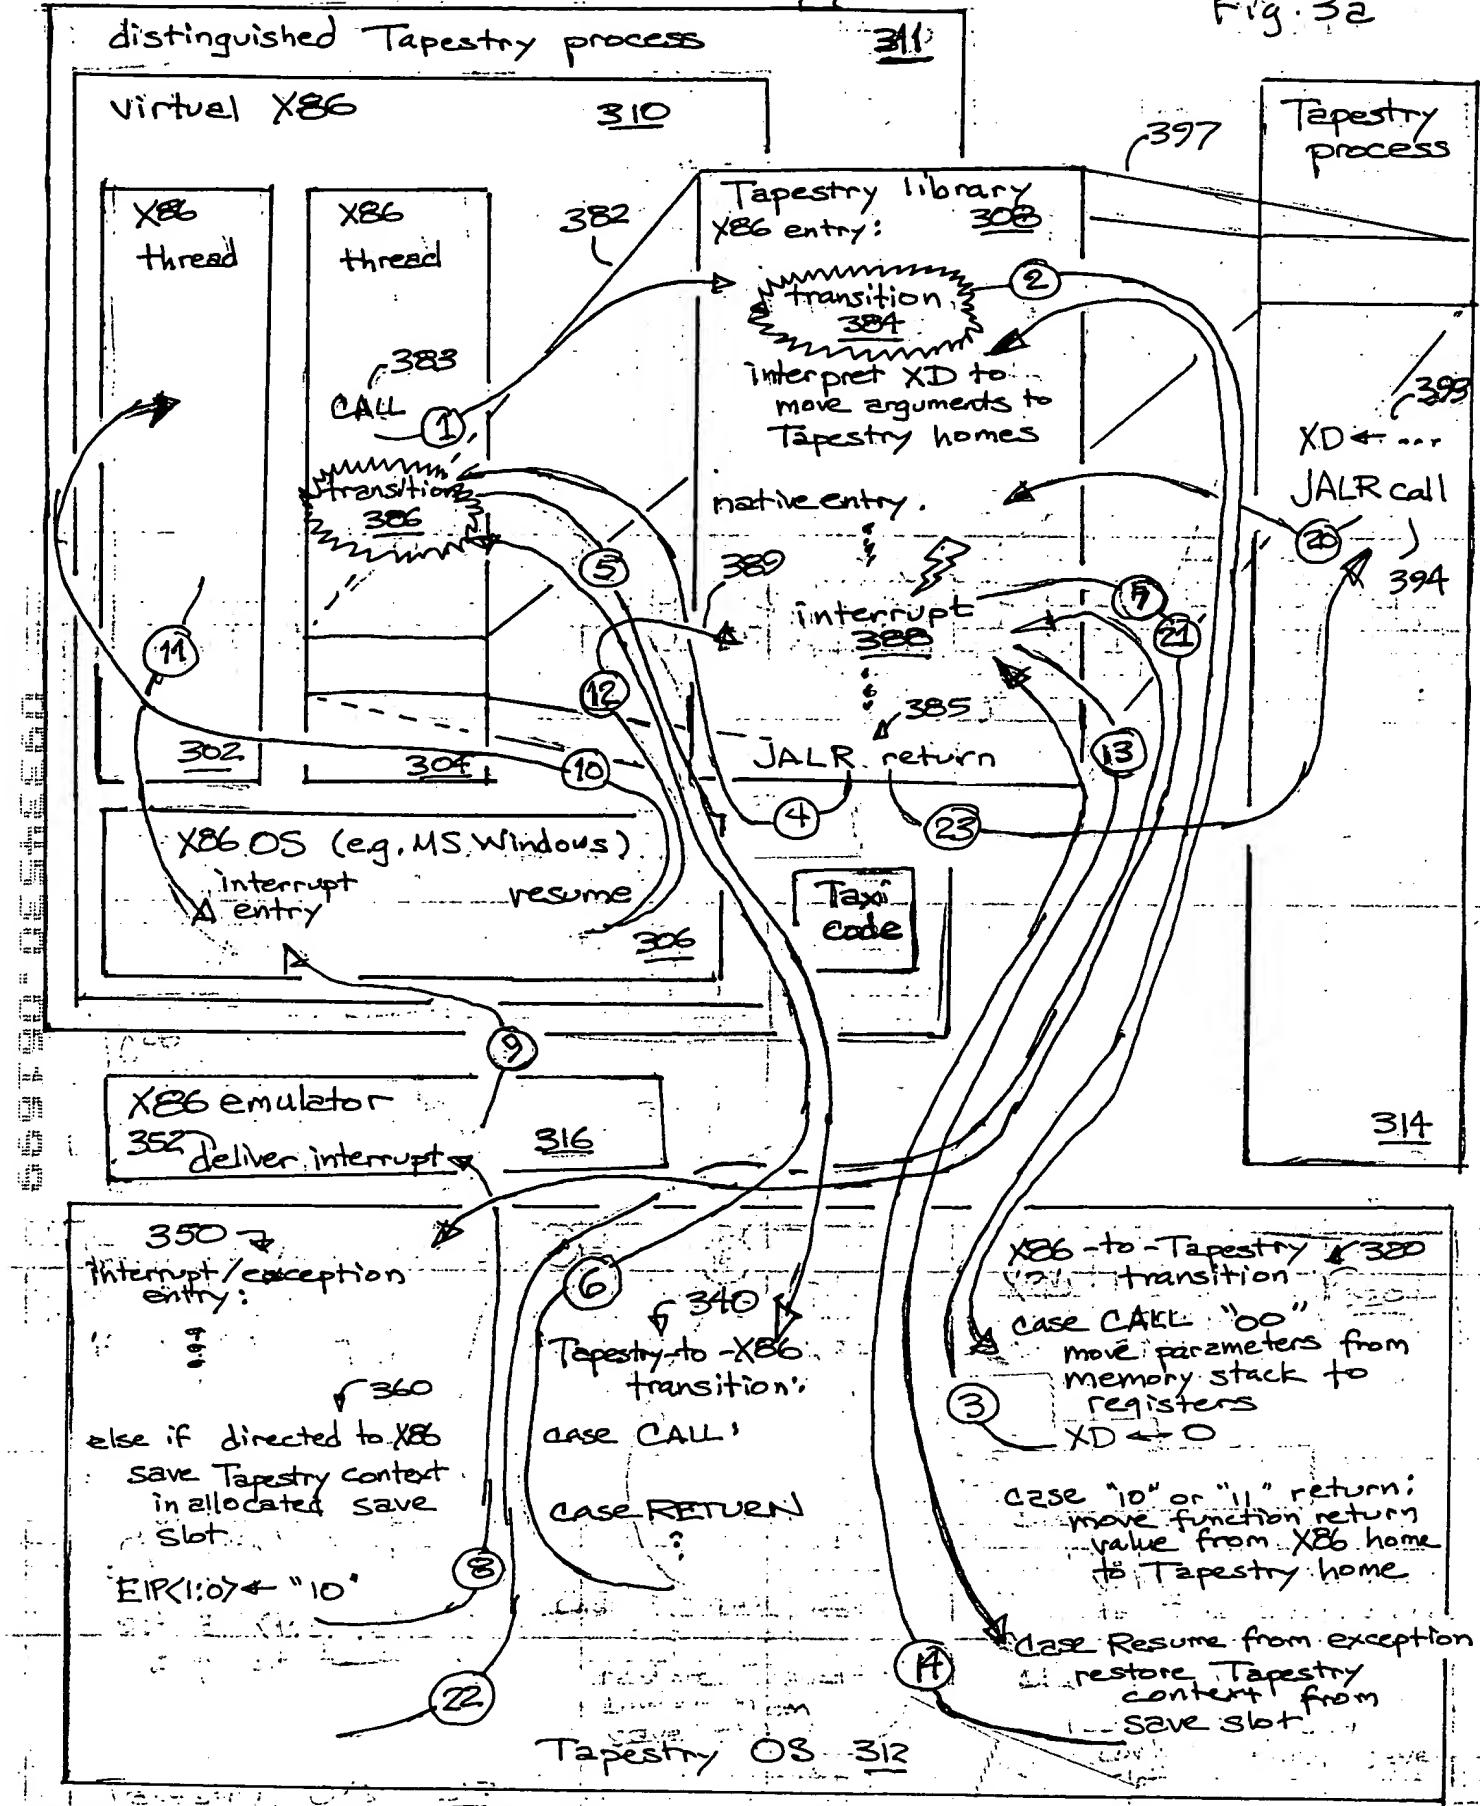
Transition (source => dest) ISA & CC property values	Handler Action
00 => 00	No transition exception
00 => 01	VECT_xxx_X86_CC exception - handler converts from native to x86 conventions
00 => 1x	VECT_xxx_X86_CC exception - handler converts from native to x86 conventions, sets up expected emulator and profiling state
01 => 00	VECT_xxx_TAP_CC exception - handler converts from x86 to native conventions
01 => 01	No transition exception
01 => 1x	VECT_X86_ISA exception [conditional based on PCW.X86_ISA_ENABLE flag] - sets up expected emulator and profiling state
1x => 00	VECT_xxx_TAP_CC exception - handler converts from x86 to native conventions
1x => 01	VECT_TAP_ISA exception [conditional based PCW.TAP_ISA_ENABLE flag] - no convention conversion necessary
1x => 10	No transition exception - [profile complete possible, probe possible]
1x => 11	No transition exception - [profile complete possible, probe NOT possible]

Fig. 2b : ISA & CC transition exception flow

	name	description	type
242	VECT_call_X86_CC	push args, return address, set up x86 state	fault on target instruction
244	VECT_jump_X86_CC	set up x86 state	fault on target instruction
246	VECT_ret_no_fp_X86_CC	return value to eax:edx, set up x86 state	fault on target instruction
248	VECT_ret_fp_X86_CC	return value to x86 fp stack, set up x86 state	fault on target instruction
250	VECT_call_TAP_CC	x86 stack args, return address to registers	fault on target instruction
252	VECT_jump_TAP_CC	x86 stack args to registers	fault on target instruction
254	VECT_ret_no_fp_TAP_CC	return value to RV0	fault on target instruction
256	VECT_ret_any_TAP_CC	return type unknown, setup RV0 and RVDP	fault on target instruction

Fig. 2c : CC transition exceptions

Fig. 32



## Flat 32-bit "Near" Address Space

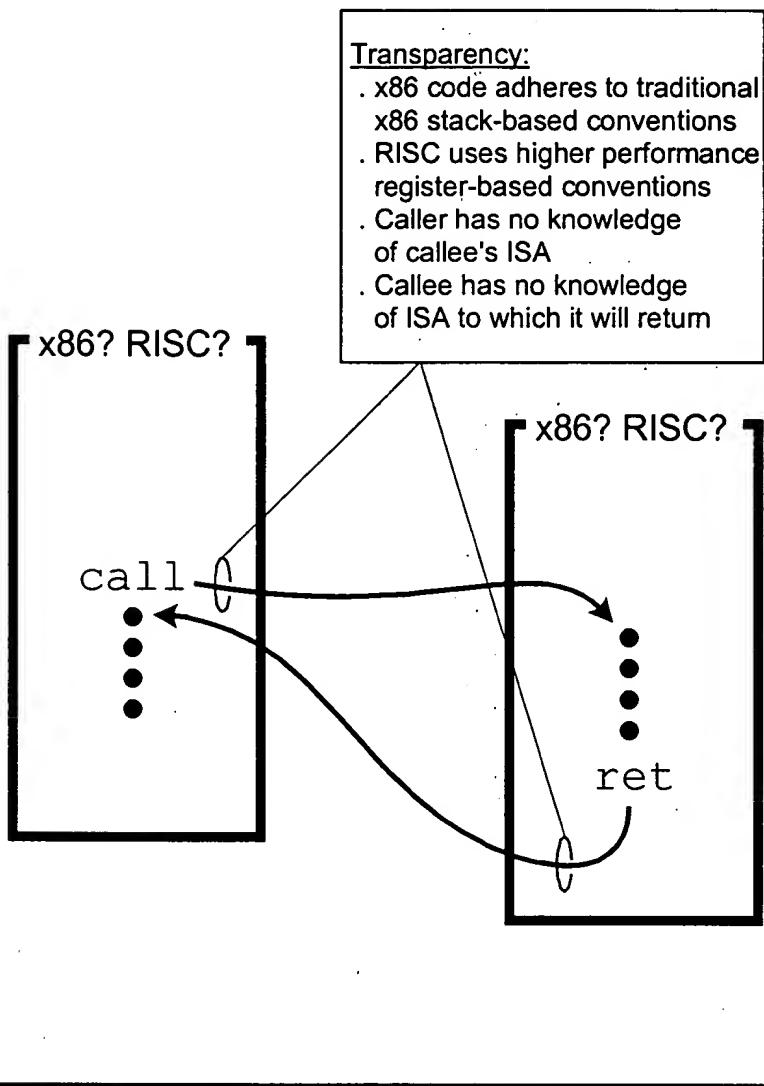


Fig. 3b

## Flat 32-bit "Near" Address Space

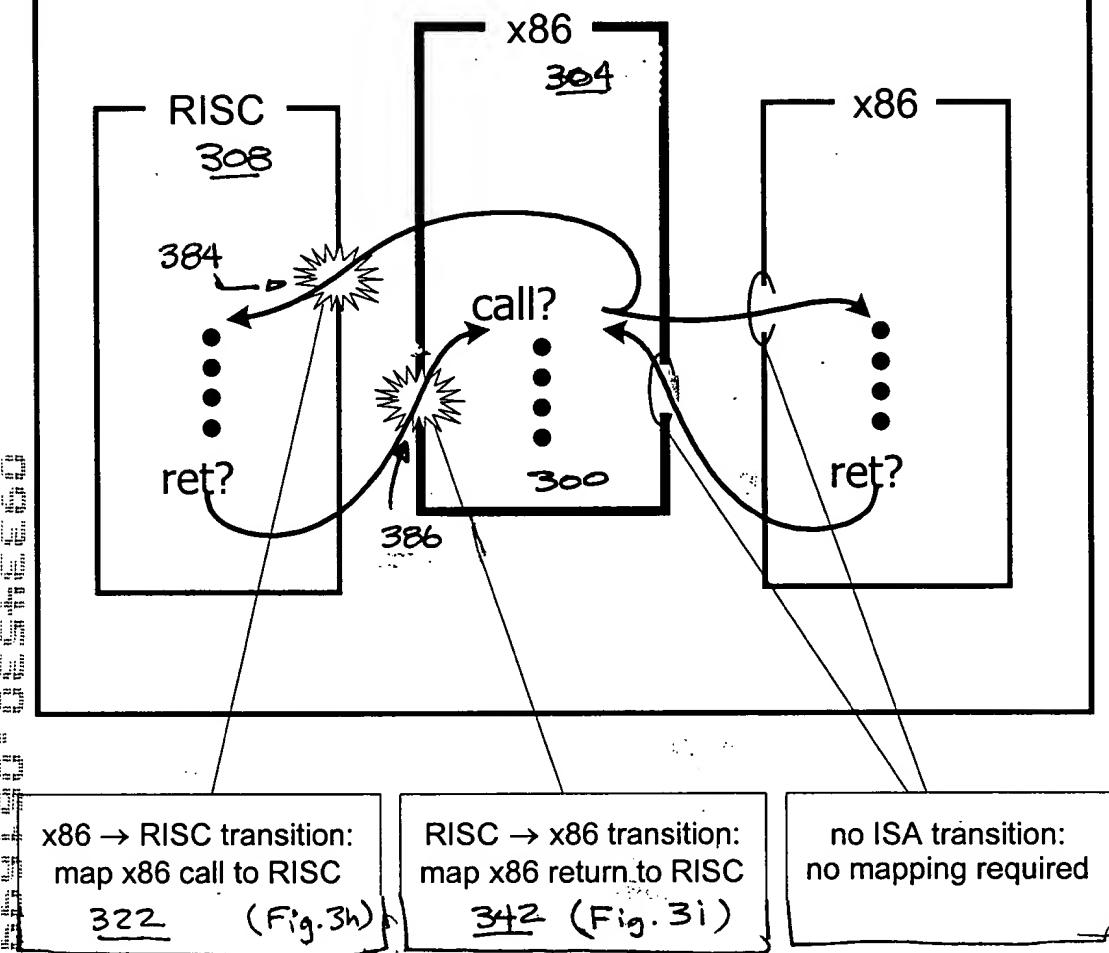


Fig. 3c

## Flat 32-bit "Near" Address Space

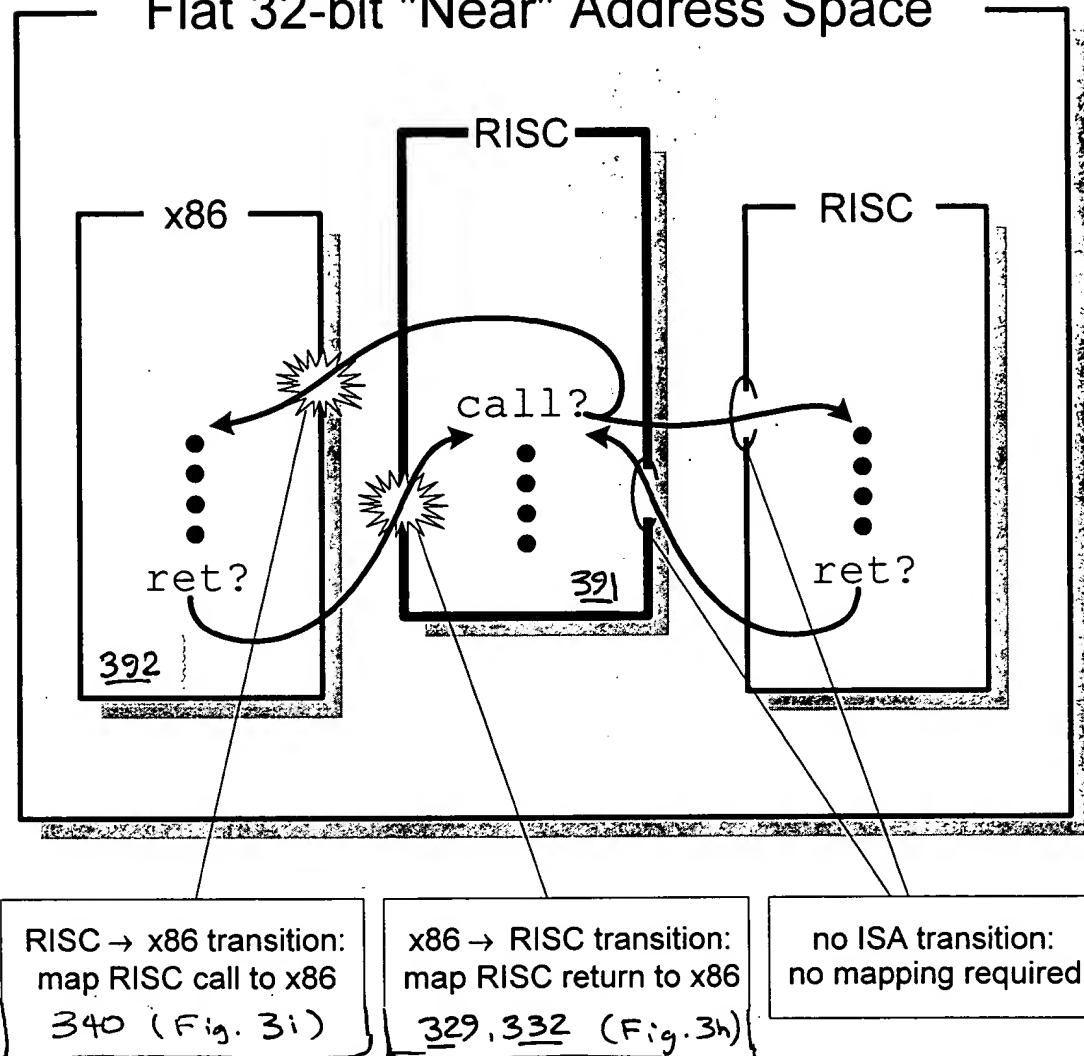
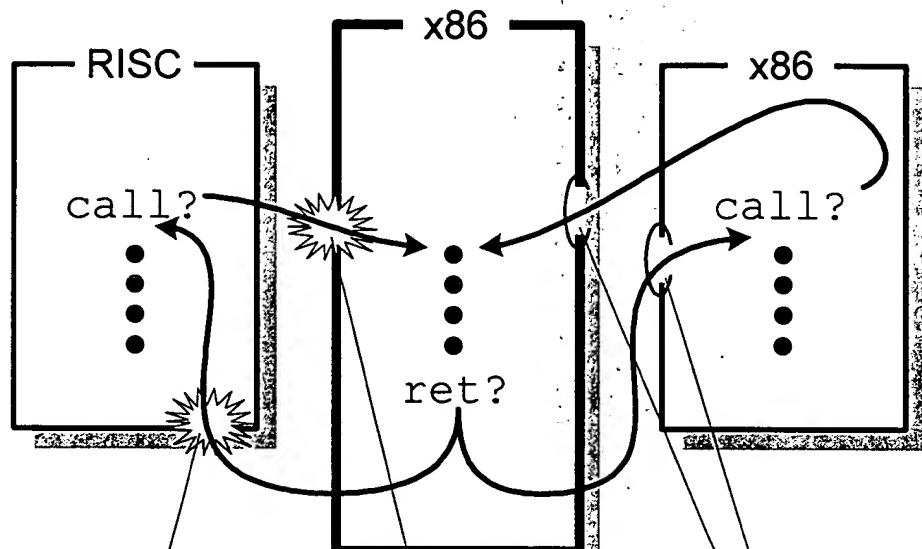


Fig. 3d

## Flat 32-bit "Near" Address Space



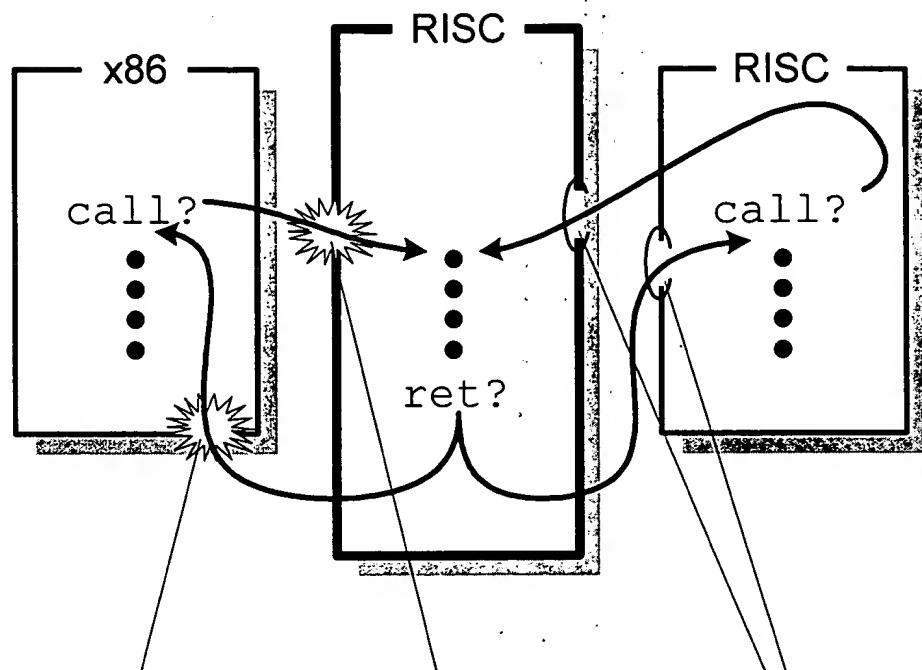
x86 → RISC transition:  
map RISC return to x86  
329, 332 (Fig. 3h)

RISC → x86 transition:  
map RISC call to x86  
343 - 348 (Fig. 3i)

no ISA transition:  
no mapping required

Fig. 3e

## Flat 32-bit "Near" Address Space



RISC → x86 transition:  
map x86 return to RISC

342 (Fig. 3i)

x86 → RISC transition:  
map x86 call to RISC

322 (Fig. 3h)

no ISA transition:  
no mapping required

Fig. 3f

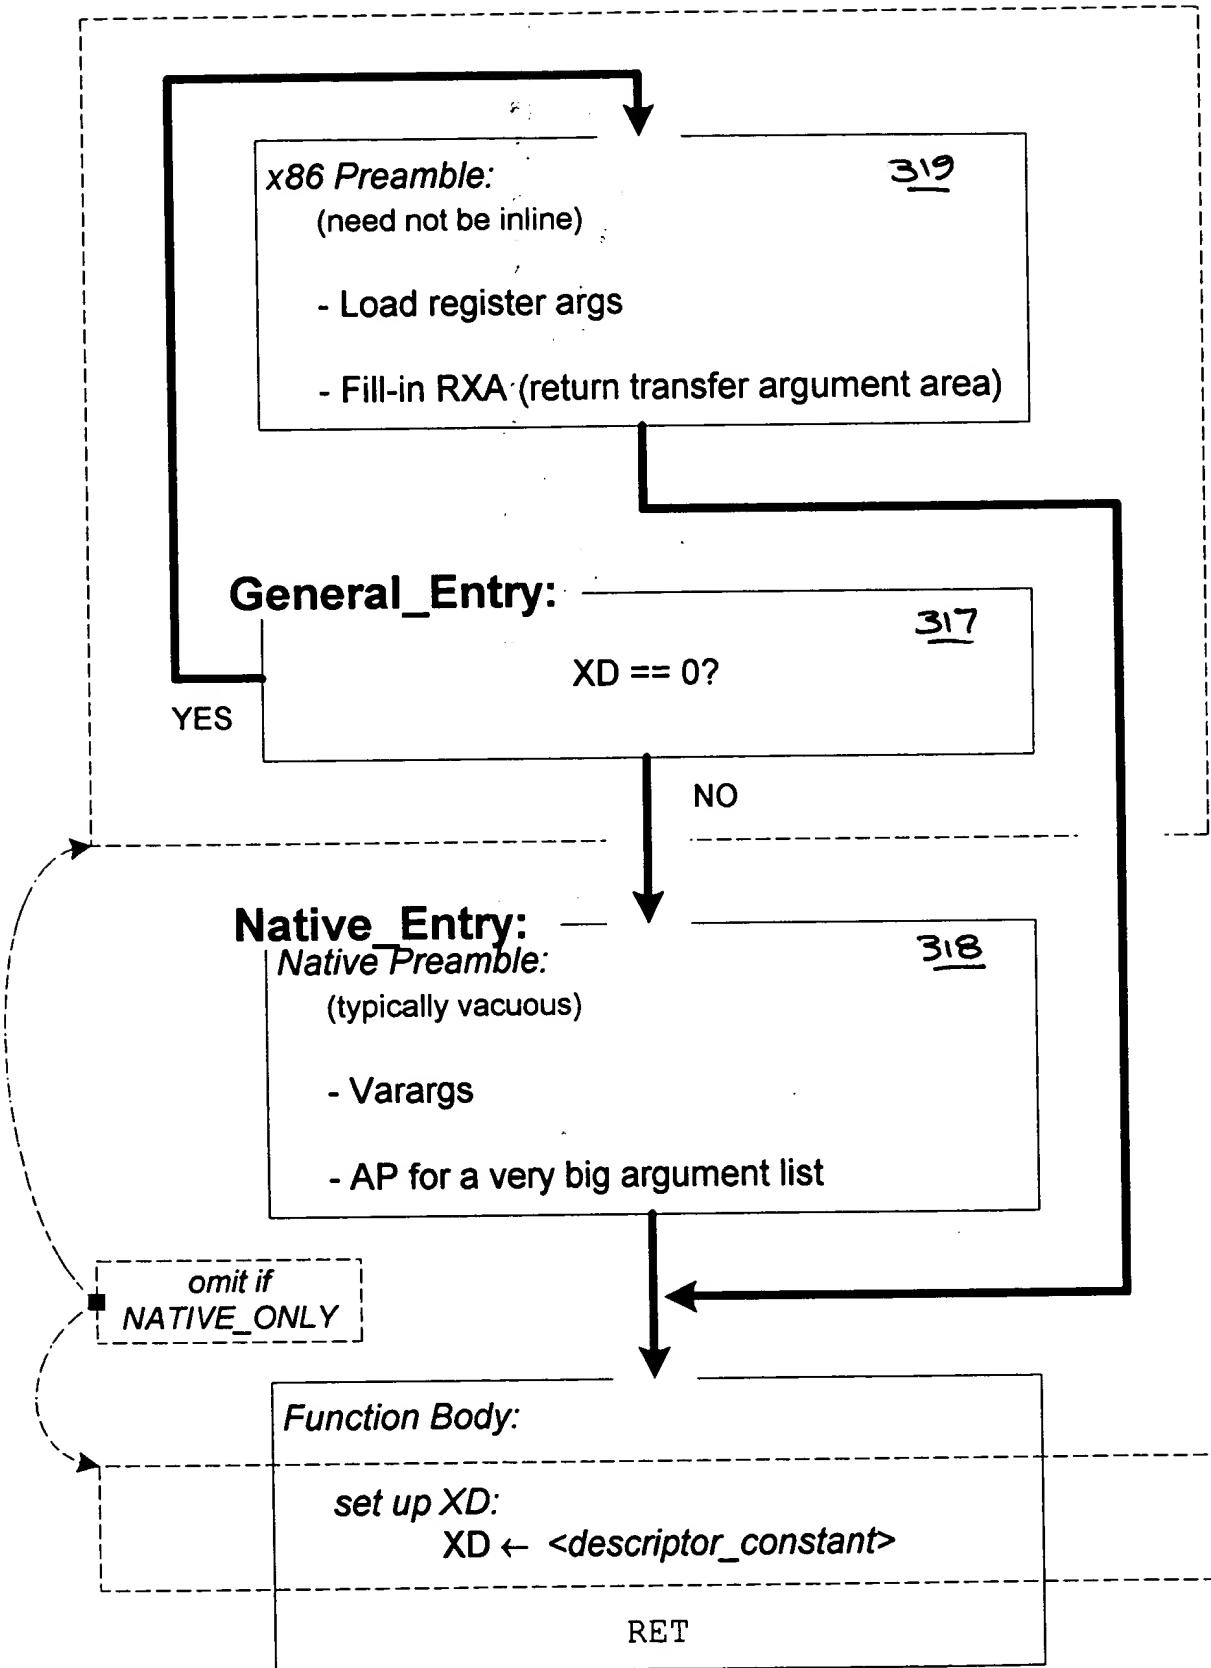


Fig. 3g

## X86-to-Tapestry transition exception handler

320

// This handler is entered under the following conditions:  
// 1. An x86 caller invokes a native function  
// 2. An x86 function returns to a native caller  
// 3. x86 software returns to or resumes an interrupted native function following  
// an external asynchronous interrupt, a processor exception, or a context switch

321

dispatch on the two least-significant bits of the destination address {  
case "00" // calling a native subprogram

// copy linkage and stack frame information and call parameters from the memory  
// stack to the analogous Tapestry registers

LR ← [SP++] // set up linkage register ~ 323

AP ← SP // address of first argument ~ 324

SP ← SP - 8 // allocate return transfer argument area ~ 326

SP ← SP & (-32) // round the stack pointer down to a 0 mode 32 boundary ~ 327

XD ← 0 // inform callee that caller uses X86 calling conventions ~ 328

322

case "01" // resuming an X86 thread suspended during execution of a native routine

if the redundant copies of the save slot number in EAX and EDX do not match or if

the redundant copies of the timestamp in EBX:ECX and ESI:EDI do not match { 371

// some form of bug or thread corruption has been detected

goto TAPESTRY\_CRASH\_SYSTEM( thread-corruption-error-code ) ~ 372

}

save the EBX:ECX timestamp in a 64-bit exception handler temporary register { 373  
(this will not be overwritten during restoration of the full native context)

use save slot number in EAX to locate actual save slot storage ~ 374

restore full entire native context (includes new values for all x86 registers) ~ 375

if save slot's timestamp does not match the saved timestamp { ~ 376

// save slot as been reallocated; save slot exhaustion has been detected

goto TAPESTRY\_CRASH\_SYSTEM( save-slot-overwritten-error-code ) ~ 377

}

free the save slot ~ 378

370

case "10" // returning from X86 callee to native caller, result already in registers

RV0<63:32> ← edx<31:00> // in case result is 64 bits ~ 333

convert the FP top-of-stack value from 80 bit X86 form to 64-bit form in RVDP ~ 334

SP ← ESI // restore SP from time of call ~ 337

332

case "11" // returning from X86 callee to native caller, load large result from memory

RV0..RV3 ← load 32 bytes from [ESI-32] // (guaranteed naturally aligned) ~ 330

SP ← ESI // restore SP from time of call ~ 337

329

}

EPC ← EPC & -4 // reset the two low-order bits to zero ~ 336

RFE ~ 338

Fig. 3h

### Tapestry-to-X86 transition exception handler

```

// This handler is entered under the following conditions: 340
// 1. a native caller invokes an x86 function
// 2. a native function returns to an x86 caller
switch on XD<3:0> { ~ 341

    XD_RET_FP:           // result type is floating point
        F0/F1 ← FINFLATE.de( RVDP ) // X86 FP results are 80 bits
        SP ← from RXA save        // discard RXA, pad, args
        FPCW ← image after FINIT & push // FP stack has 1 entry
        goto EXIT

    XD_RET_WRITEBACK:      // store result to @RVA, leave RVA in eax
        RVA ← from RXA save    // address of result area
        copy decode(XD<8:4>) bytes from RV0..RV3 to [RVA]
        eax ← RVA              // X86 expects RVA in eax
        SP ← from RXA save    // discard RXA, pad, args
        FPCW ← image after FINIT // FP stack is empty
        goto EXIT

    XD_RET_SCALAR:         // result in eax:eda
        edx<31:00> ← eax<63:32> // in case result is 64 bits
        SP ← from RXA save    // discard RXA, pad, args
        FPCW ← image after FINIT // FP stack is empty
        goto EXIT

    XD_CALL_HIDDEN_TEMP: // allocate 32 byte aligned hidden temp
        esi ← SP              // stack cut back on return ~ 343
        SP ← SP - 32           // allocate max size temp
        RVA ← SP              // RVA consumed later by RR } 344
        LR<1:0> ← "11"        // flag address for return & reload ~ 345
        goto CALL_COMMON

    default:                // remaining XD_CALL_xxx encodings
        esi ← SP              // stack cut back on return ~ 343
        LR<1:0> ← "10"        // flag address for return ~ 346

CALL_COMMON:
    interpret XD to push and/or reposition args ~ 347
    [-SP] ← LR              // push LR as return address } 348

EXIT:
    setup emulator context and profiling ring buffer pointer
}

RFE ~ 349                // to original target
}

```

Fig. 3i

## interrupt/exception handler of Tapestry operating system:

// Control vectors here when a synchronous exception or asynchronous interrupt is to be  
// exported to / manifested in an x86 machine.

// The interrupt is directed to something within the virtual X86, and thus there is a possibility  
// that the X86 operating system will context switch. So we need to distinguish two cases:  
// either the running process has only X86 state that is relevant to save, or  
// there is extended state that must be saved and associated with the current machine context  
// (e.g., extended state in a Tapestry library call in behalf of a process managed by X86 OS)  
if execution was interrupted in the converter – EPC.ISA == X86 {

// no dependence on extended/native state possible hence no need to save any  
goto EM86\_Deliver\_Interrupt( interrupt-byte ) } 351

} else if EPC.Taxi\_Active {

// A Taxi translated version of some X86 code was running. Taxi will rollback to an  
// x86 instruction boundary. Then, if the rollback was induced by an asynchronous external  
// interrupt Taxi will deliver the appropriate x86 interrupt. Else, the rollback was induced  
// by a synchronous event so Taxi will resume execution in the converter, retriggering the  
// exception but this time will EPC.ISA == X86  
goto TAXi\_Rollback( asynchronous-flag, interrupt-byte ) } 353

} else if EPC.EM86 {

// The emulator has been interrupted. In theory the emulator is coded to allow for such  
// conditions and permits re-entry during long running routines (e.g. far call through a gate)  
// to deliver external interrupts  
goto EM86\_Deliver\_Interrupt( interrupt-byte ) } 354

else {

// This is the most difficult case – the machine was executing native Tapestry code on  
// behalf of an X86 thread. The X86 operating system may context switch. We must save  
// all native state and be able to locate it again when the x86 thread is resumed.

361  
allocate a free save slot; if unavailable free the save slot with oldest timestamp and try again  
save the entire native state (both the X86 and the extended state) } 362  
save the X86 EIP in the save slot  
overwrite the two low-order bits of EPC with "01" (will become X86 interrupt EIP) ~ 363  
store the 64-bit timestamp in the save slot, in the X86 EBX:ECX register pair (and, } 364  
for further security, store a redundant copy in the X86 ESI:EDI register pair)  
store the a number of the allocated save slot in the X86 EAX register (and, again for } 365  
further security, store a redundant copy in the X86 EDX register)  
goto EM86\_Deliver\_Interrupt( interrupt-byte ) ~ 369 } 360

Fig. 3j

```

typedef struct {
    save_slot_t * newer;           // pointer to next-most-recently-allocated save slot } 379c
    save_slot_t * older;           // pointer to next-older save slot
    unsigned int64 epc;            // saved exception PC/IP
    unsigned int64 pcw;            // saved exception PCW (program control word)
    unsigned int64 registers[63];  // save the 63 writeable general registers } 356
    ...                           // other words of Tapestry context
    timestamp_t timestamp;         // timestamp to detect buffer overrun ~ 358
    int save_slot_ID;             // ID number of the save slot ~ 357
    boolean save_slot_is_full;     // full / empty flag ~ 359
} save_slot_t;

save_slot_t * save_slot_head;    // pointer to the head of the queue ~ 379a
save_slot_t * save_slot_tail;   // pointer to the tail of the queue ~ 379b

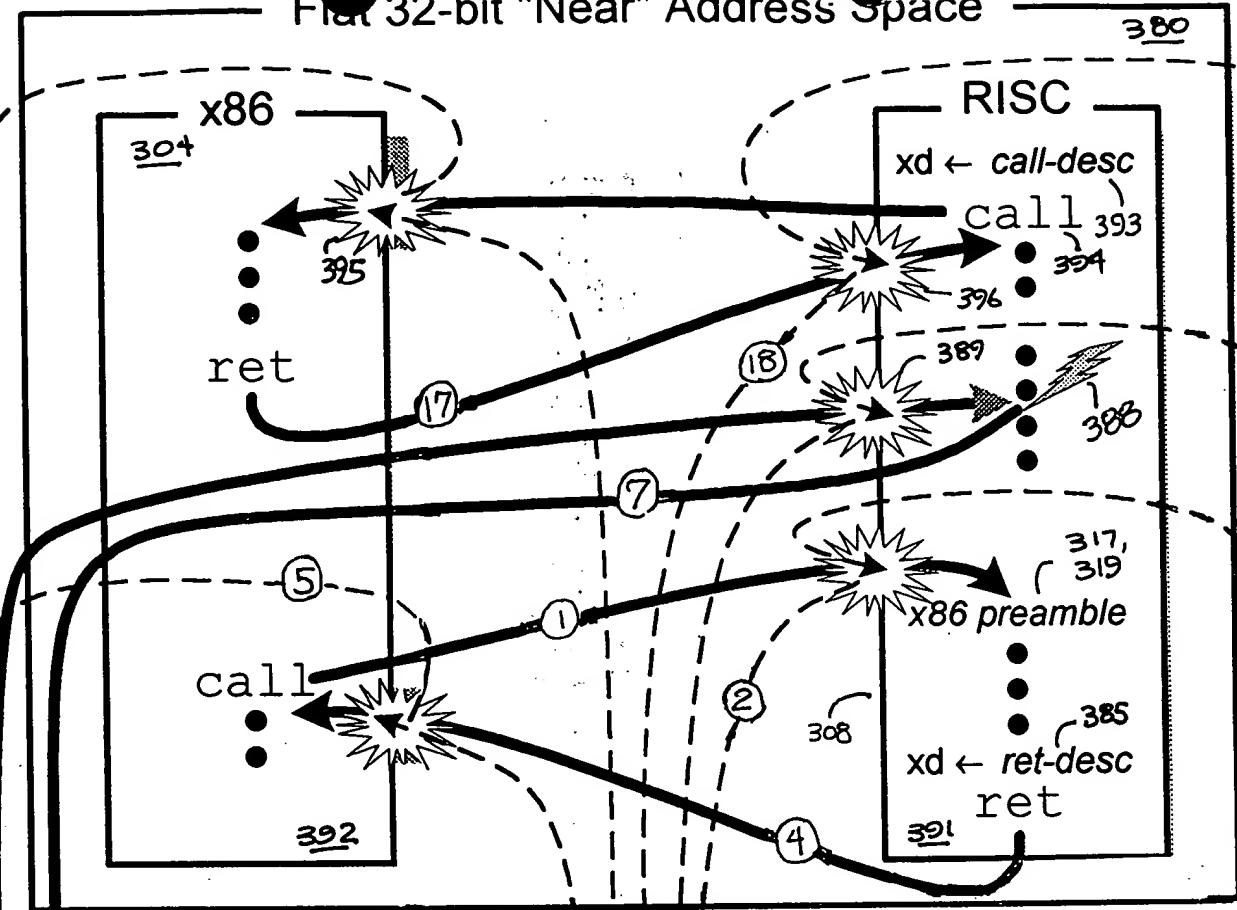
```

**System initialization**  
reserve several pages of unpaged memory for save slots

**Fig. 3k**

Fig. 31

## Flat 32-bit "Near" Address Space



## Prepare x86 excep. or int.

- Alloc free or oldest save slot
- Store timestamp & full state
- x86 regs ← save slot ID, TS
- EPC<1:0> ← 01

## Handler: x86 to RISC

EPC<1:0> = 00: 322

- LR ← [SP]
- SP ← SP + 4
- AP ← SP
- SP ← SP - 8 // ret area
- SP ← SP & (-32)
- XD ← 0

EPC<1:0> = 01: 370

- x86 regs points to save slot
- Using TS verify no overwrite
- Restore full state
- Free save slot
- EPC<1:0> ← 00

EPC<1:0> = 1x: 329 332

- Reformat / reposition the function result per EPC<0>
- SP ← ESI
- EPC<1:0> ← 00

12

x86 SW

306, 316, 302

## Handler: RISC to x86

XD contains return-descriptor:

- Interpret XD: 342
- Reformat / reposition result
- Load FPCW
- SP ← [SP] // pop RA & args

XD contains call-descriptor:

- ESI ← SP
- Interpret XD, reposition args
- LR<1:0> ← 1x per XD
- Push LR as RA (ret addr)

Fig. 31

## Flat 32-bit "Near" Address Space

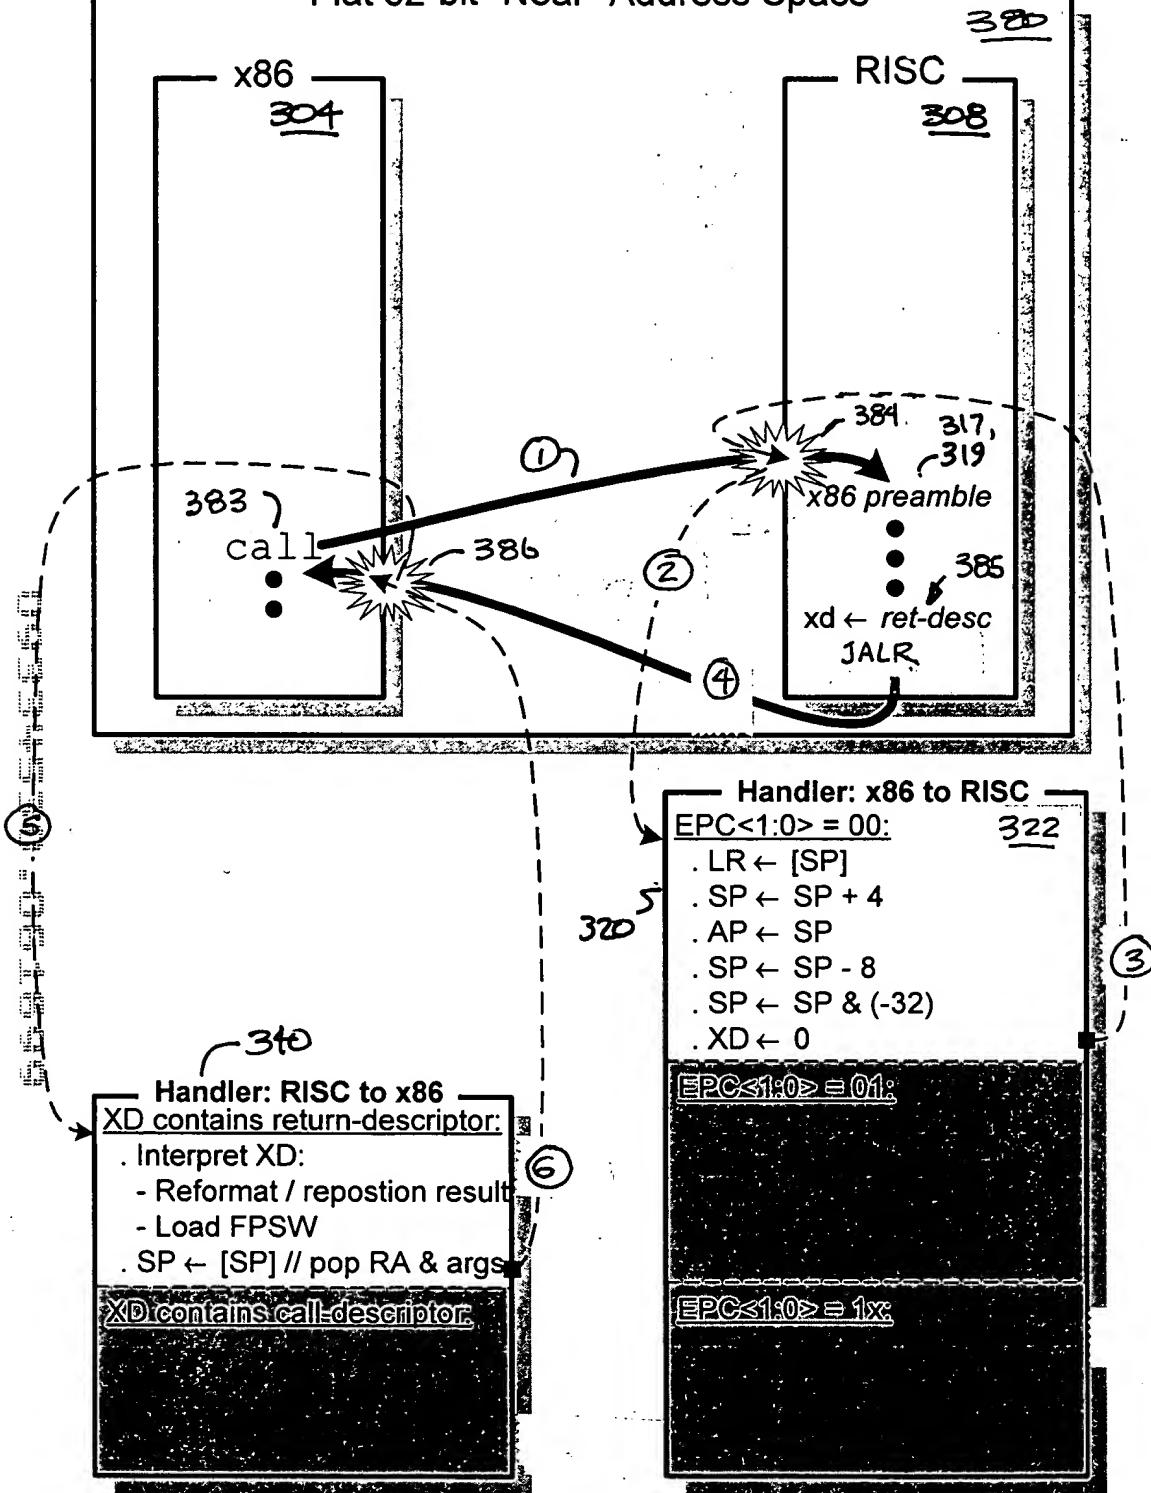


Fig. 3m

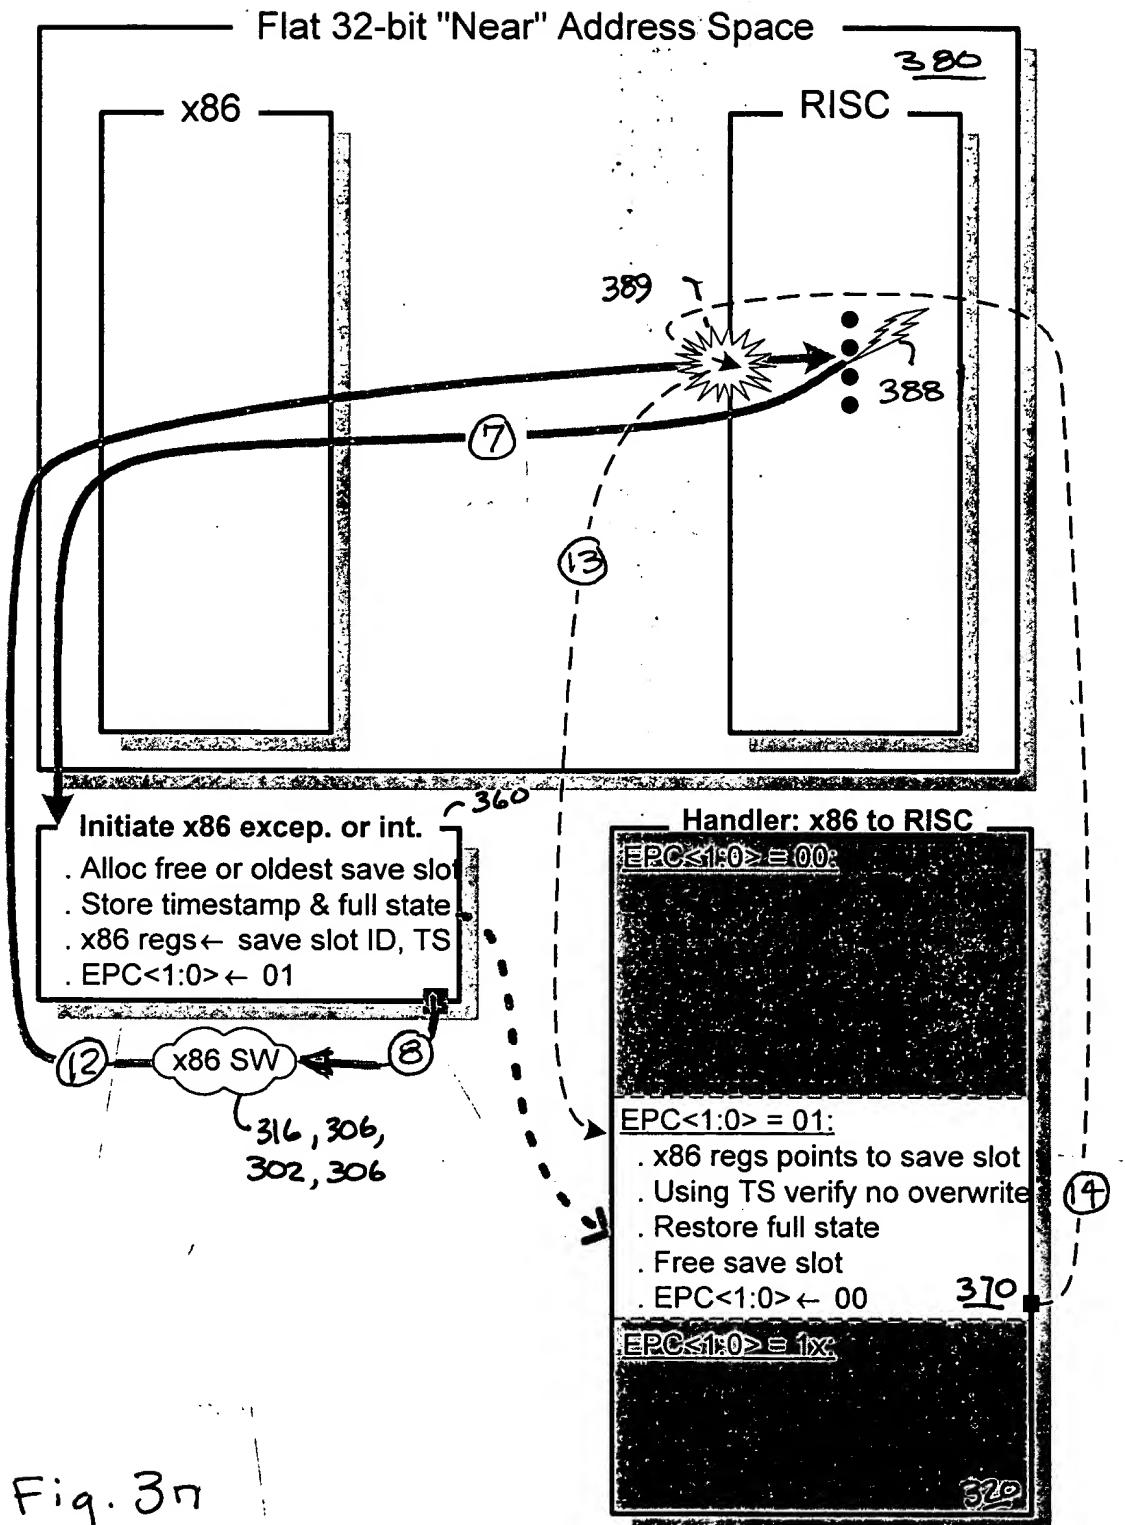


Fig. 37

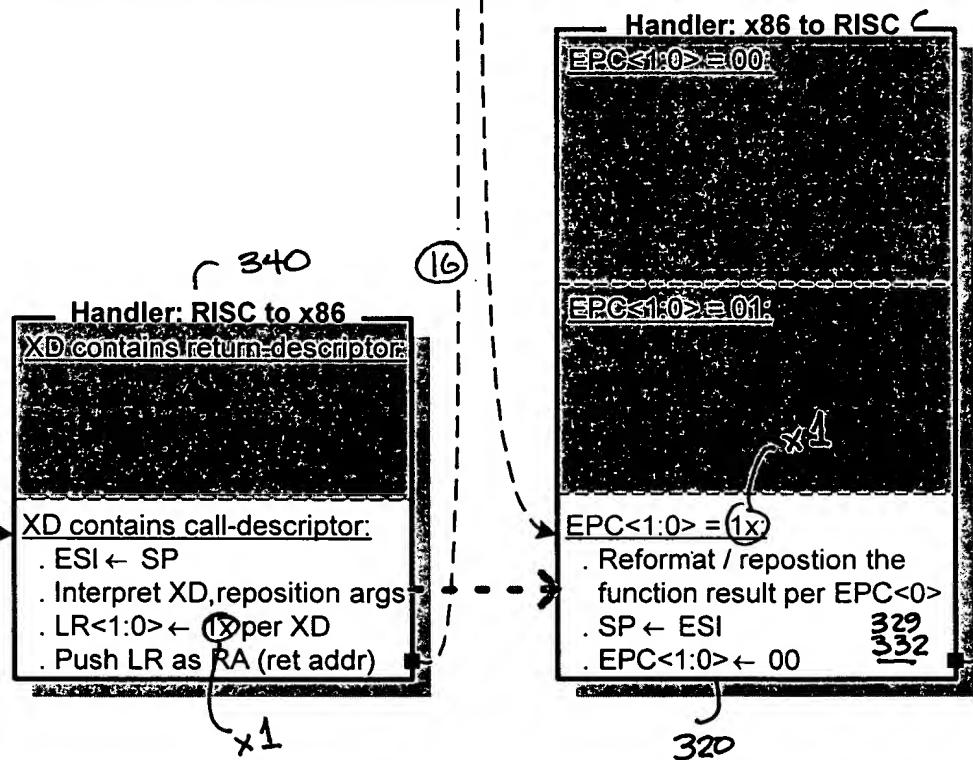
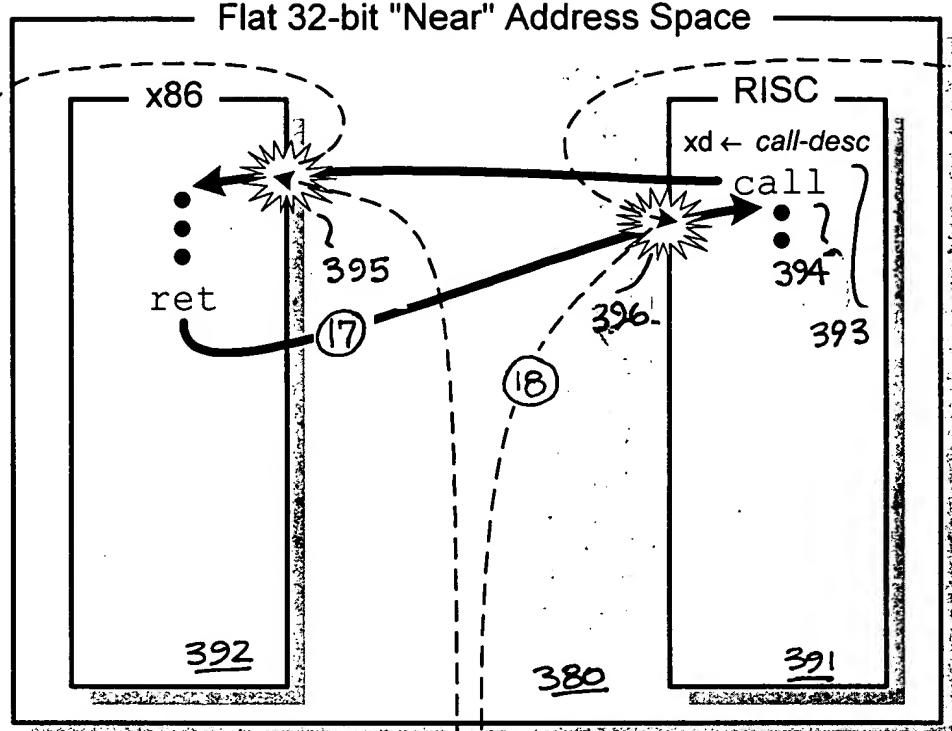
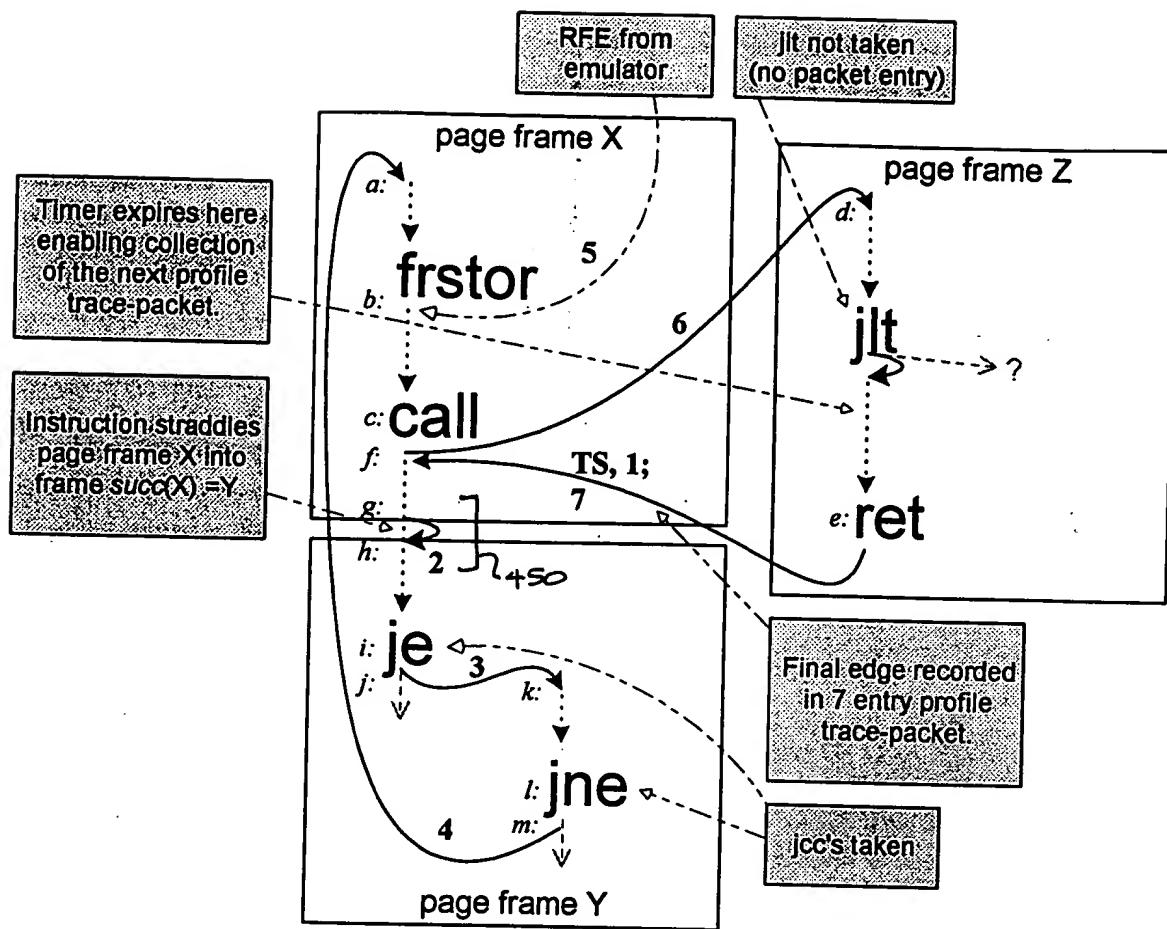


Fig. 30



### 7 entry trace packet

420 {

Entry	Event Code	Done Addr	Next Addr
64 bit time stamp			
1	ret	x86 context	phys X:f
2	new page	phys Y:g	phys Y:h
3	jcc forward	phys Y:i	phys Y:k
4	jnz backward	phys Y:l	phys X:a
5	seq; env change	x86 context	phys X:b
6	ip-rel near call	phys X:c	phys Z:d
7	near ret	phys Z:e	phys X:f

~ 450  
~ 440, 454  
~ 440  
~ 440  
~ 430  
~ 440  
~ 440

Fig.42

Source	Code	Event	Reuse event code	Profileable event	Initiate packet	Probeable event	Probe event bit - TLB probe attribute or Emulator probe
	402						
RFE (Context_at_Point entry)	0.0000	Default (x86 transparent) event, reuse all converter values	yes	no	no	no	
	0.0001	Simple x86 instruction completion (reuse event code)	yes	no	no	no	
	0.0010	Probe exception failed	yes	no	no	no	reuse event code
	0.0011	Probe exception failed, reload probe timer	yes	no	no	no	
	0.0100	flush event	no	no	no	no	-
	0.0101	Sequential; execution environment changed - force event	no	yes	no	no	-
	0.0110	Far RET	no	yes	yes	no	-
	0.0111	IRET	no	yes	no	no	-
	0.1000	Far CALL	no	yes	yes	yes	Far call
	0.1001	Far JMP	no	yes	yes	no	-
	0.1010	Special; emulator execution, supply extra instruction data <sup>a</sup>	no	yes	no	no	-
	0.1011	Abort profile collection	no	no	no	no	-
	0.1100	x86 synchronous/asynchronous interrupt w/probe (GRP 0)	no	yes	yes	yes	Emulator probe
	0.1101	x86 synchronous/asynchronous interrupt (GRP 0)	no	yes	yes	no	-
	0.1110	x86 synchronous/asynchronous interrupt w/probe (GRP 1)	no	yes	yes	yes	Emulator probe
	0.1111	x86 synchronous/asynchronous interrupt (GRP 1)	no	yes	yes	no	-
Converter (Near_Edge entry)	1.0000	IP-relative JNZ forward (opcode: 75, 0F 85)	no	yes	yes	no	-
	1.0001	IP-relative JNZ backward (opcode: 75, 0F 85)	no	yes	yes	yes	Jnz
	1.0010	IP-relative conditional jump forward - (Jcc, Jcxz, loop)	no	yes	yes	no	-
	1.0011	IP-relative conditional jump backward - (Jcc, Jcxz, loop)	no	yes	yes	yes	Cond jump
	1.0100	IP-relative, near JMP forward (opcode: E9, EB)	no	yes	yes	no	-
	1.0101	IP-relative, near JMP backward (opcode: E9, EB)	no	yes	yes	yes	Near jump
	1.0110	RET/ RET imm16 (opcode C3, C2 /w)	no	yes	yes	no	-
	1.0111	IP-relative, near CALL (opcode: E8)	no	yes	yes	yes	Near call
	1.1000	REPE/REPNE CMPS/SCAS (opcode: A6, A7, AE, AF)	no	yes	no	no	-
	1.1001	REP MOVS/STOS/LDOS (opcode: A4, A5, AA, AB, AC, AD)	no	yes	no	no	-
	1.1010	Indirect near JMP (opcode: FF /4)	no	yes	yes	no	-
	1.1011	Indirect near CALL (opcode: FF /2)	no	yes	yes	yes	Near call
	1.1100	load from I/O memory (TLB.asi != 0) { not used in T1 }	no	yes	no	no	-
	1.1101	available for expansion	no	no	no	no	
	1.1110	Default converter event; sequential	406	no	no	no	-
	1.1111	New page (instruction ends on last byte of a page frame or straddles across a page frame boundary)	408	no	yes	no	-

a. Used by emulator for new x86 opcodes. Extra information supplied in *Taxi\_Control.special\_opcode* bits.

Fig. 4b

432 → 433 → 434 → 435 → x86 FP Stack state

431 → sizes → modes → 434

432 → 433 → 434 → 435

Event code	436	Next: First Byte Page Frame # 438	Done: First Byte Offset 435
0	0	0	0
1	0	1	1
2	1	0	0
3	0	9	9
4	8	7	8
5	7	6	7
6	6	5	6
7	5	5	5
8	5	5	5
9	5	5	5

Event code	437	Next: First Byte Page Frame # 438	Done: First Byte Offset 435
0	0	0	0
1	0	1	1
2	1	0	0
3	0	9	9
4	8	7	8
5	7	6	7
6	6	5	6
7	5	5	5
8	5	5	5
9	5	5	5

430 → 431 → 432 → 433 → 434 → 435 → x86 FP Stack state

Fig. 4c) Context\_At\_Point profile trace-packet entry

441

Done: Length (always > 0)	Done: Last Byte Page Frame # 444	Done: First Byte Offset 445
0	6	6
1	6	6
2	5	5
3	5	5
4	5	5
5	5	5
6	5	5
7	5	5
8	5	5
9	5	5

Converter event	446	Next: First Byte Page Frame # 448	Next: First Byte Offset 449
0	3	2	2
1	0	8	8
2	8	7	7
3	5	4	4
4	3	2	2
5	2	1	1
6	5	4	4
7	6	5	5
8	7	6	6
9	8	7	7

440 →

Fig. 4d) Near\_Edge profile trace-packet entry

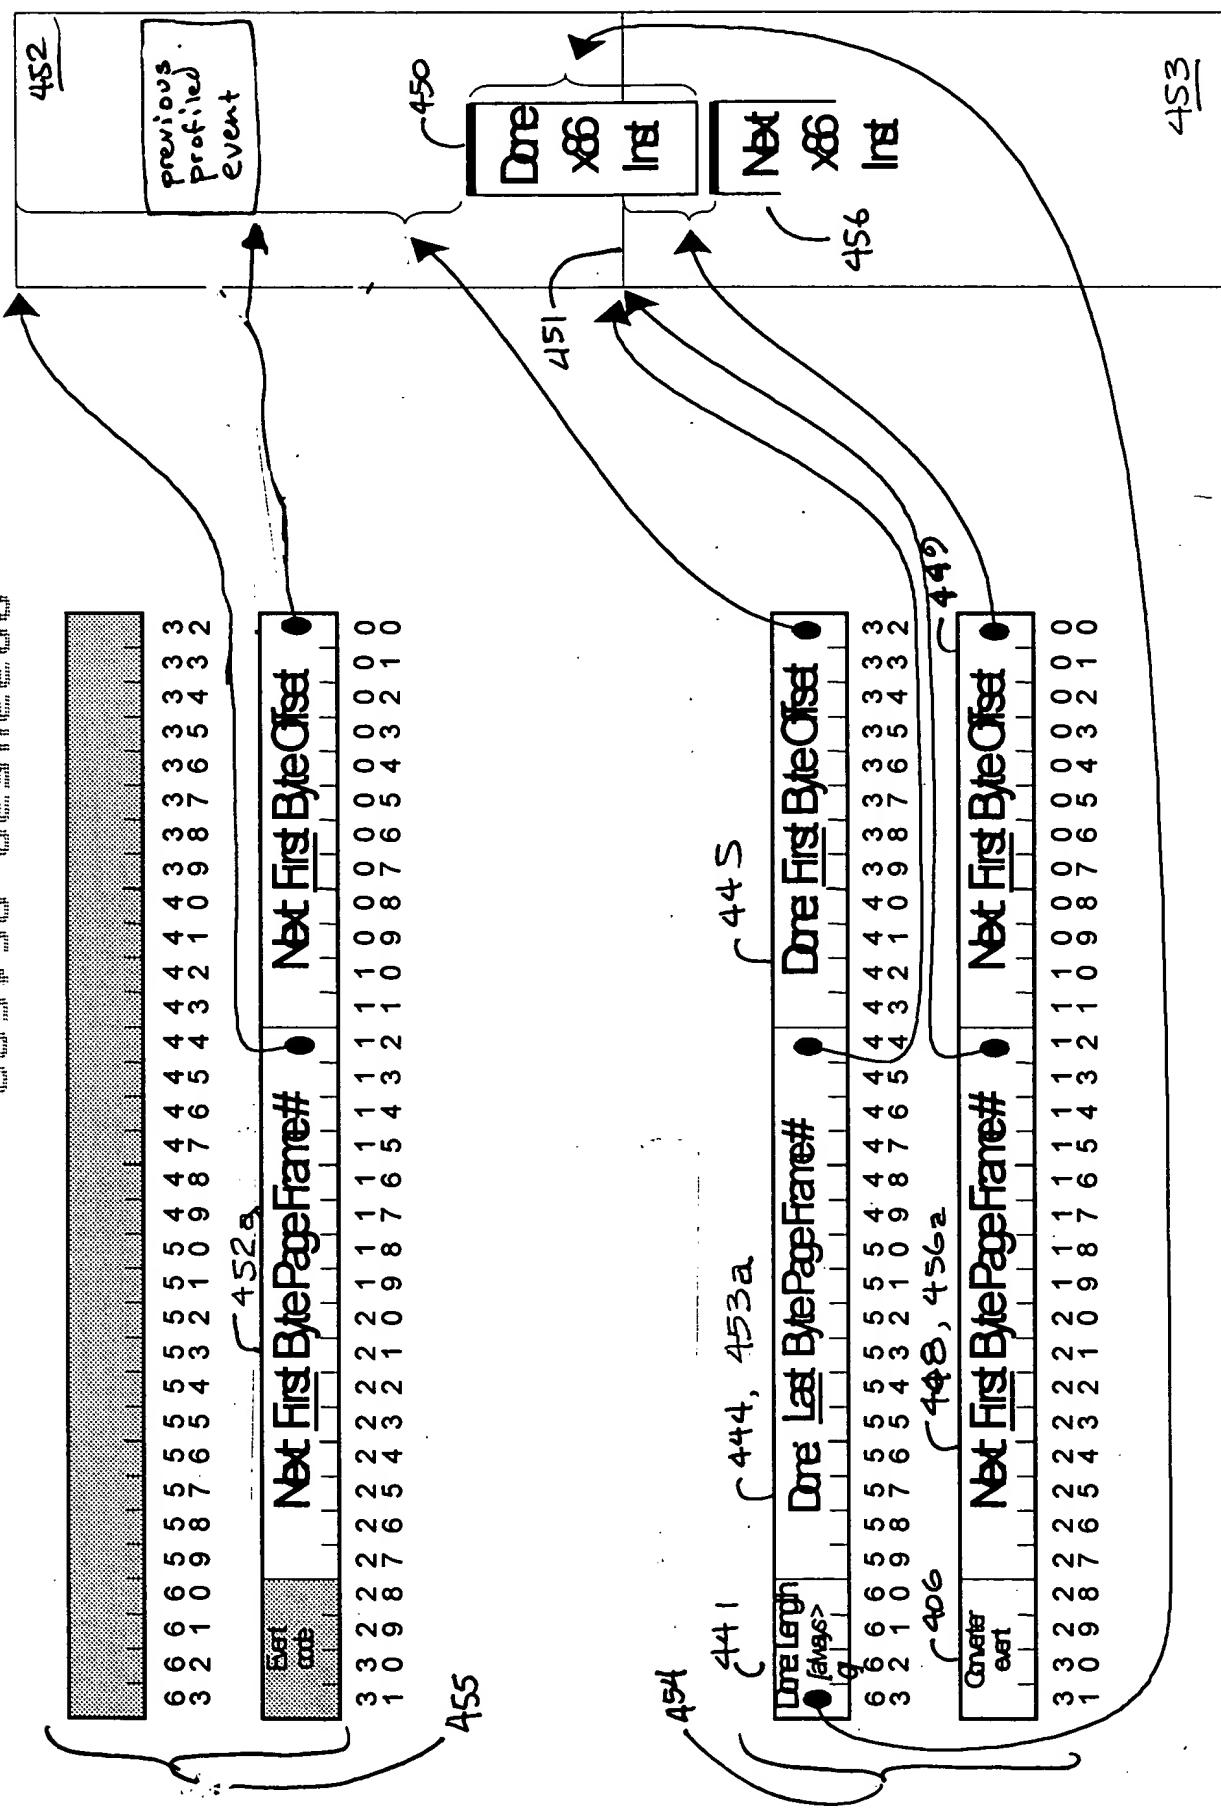
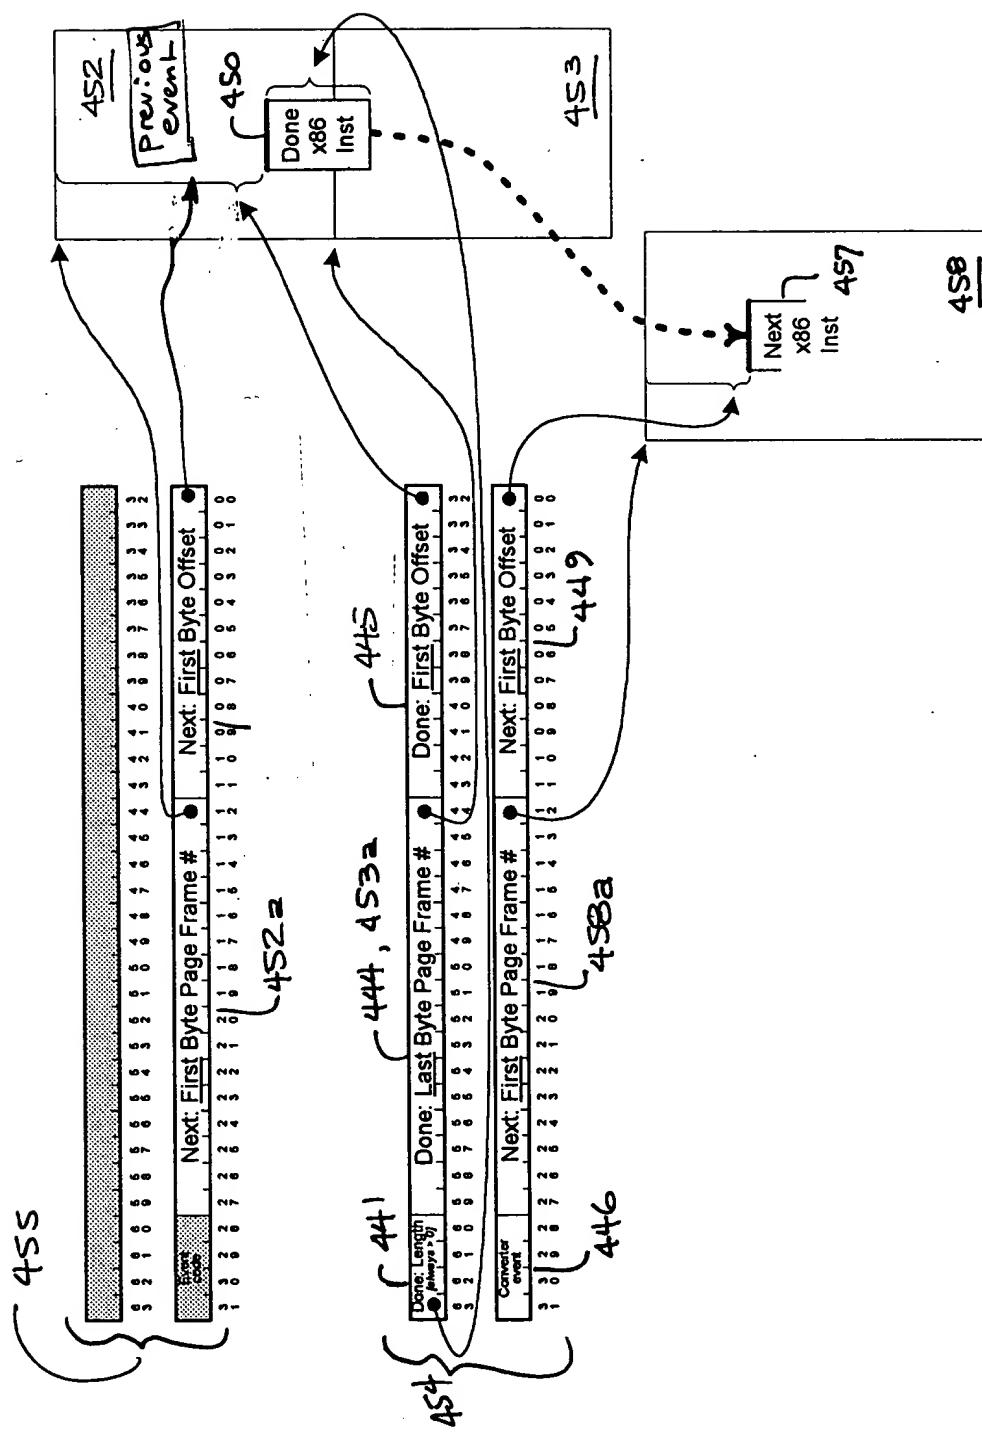


Fig 4e

Fig. 4f



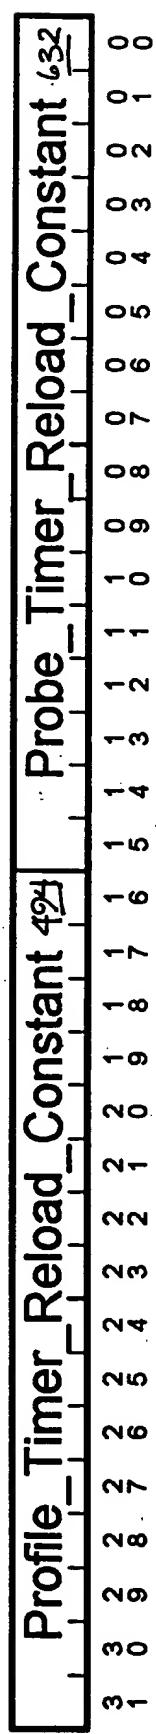
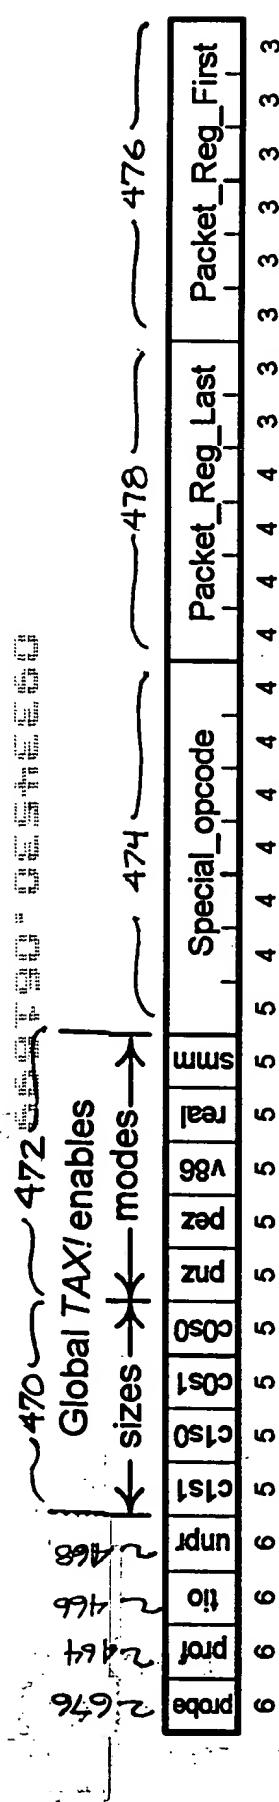


Fig. 48. Taxi\_Control processor register 460

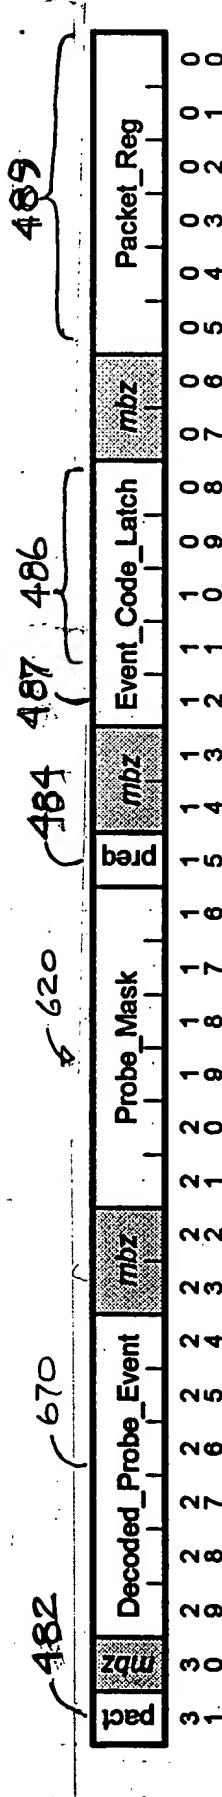


Fig. 4h Taxi\_State processor register

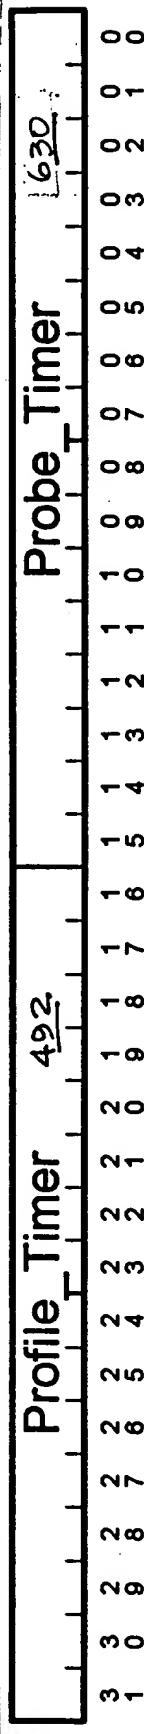
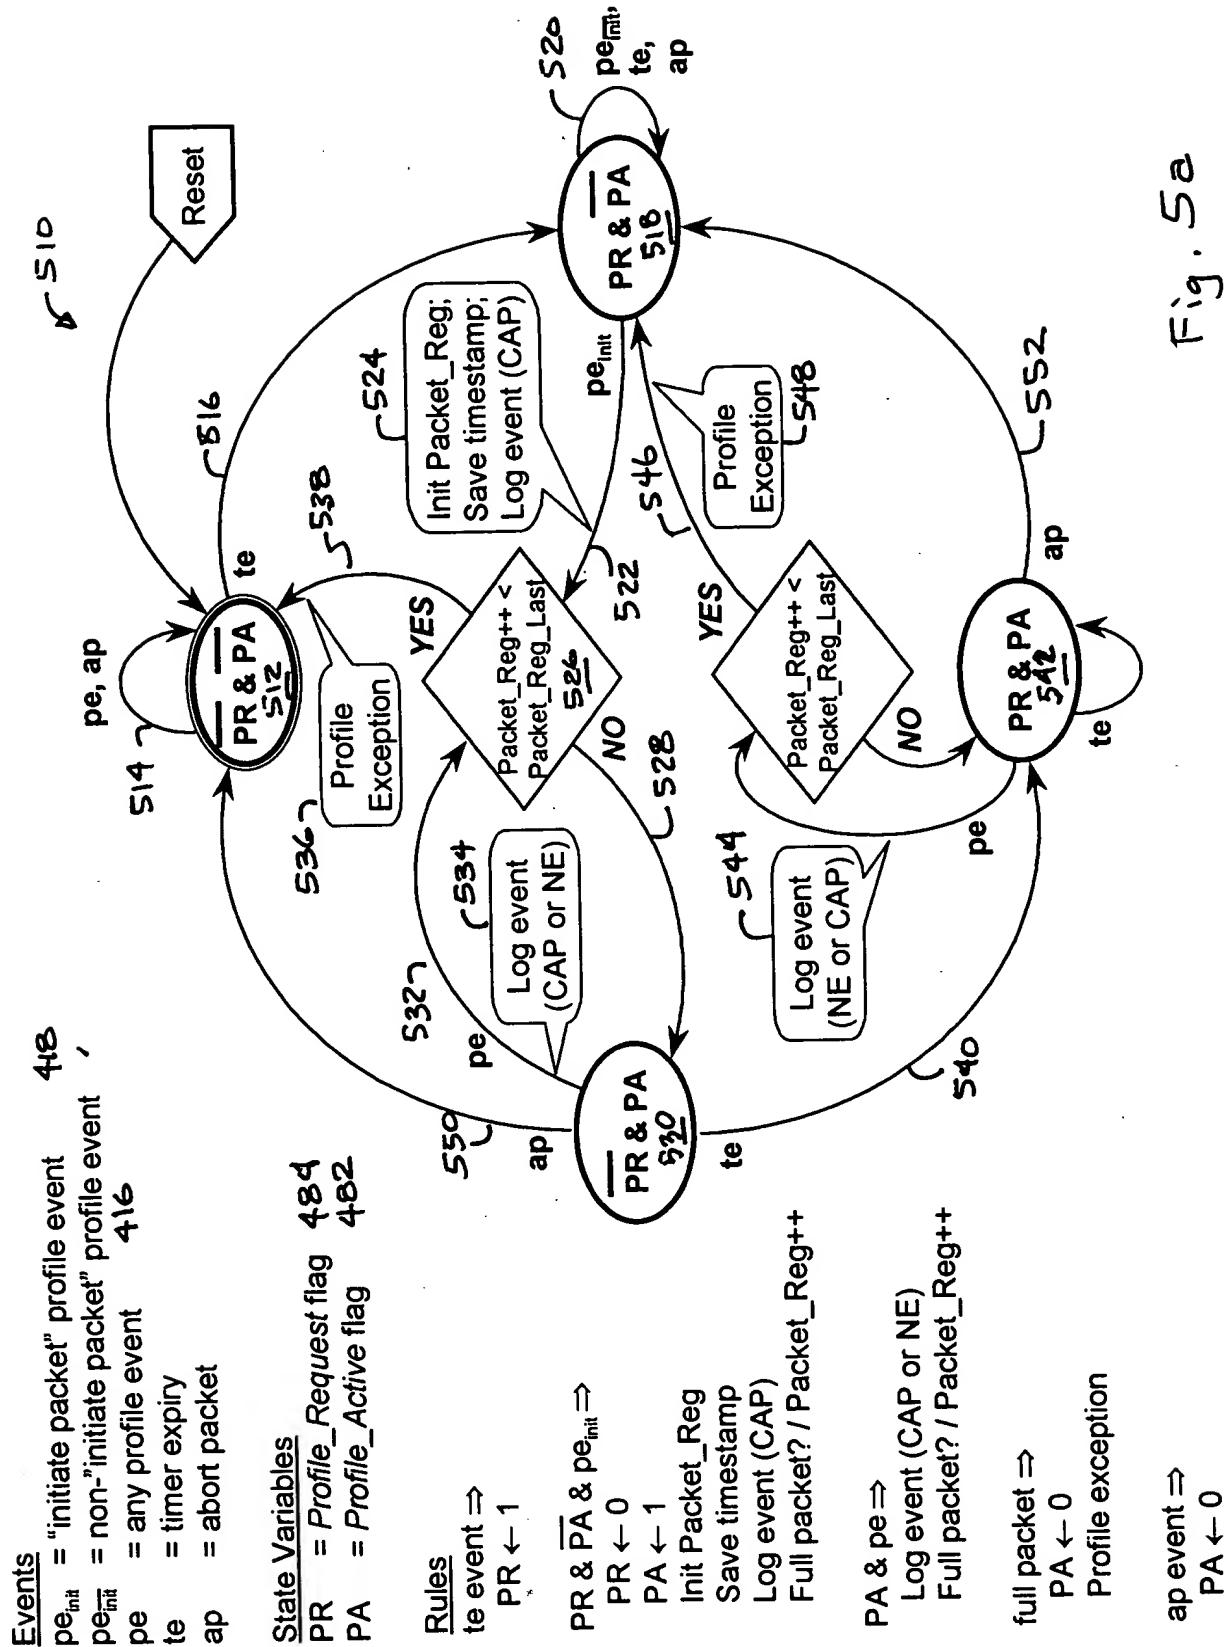


Fig. 4: *Taxi\_Timers processor register* 490



## **taxi profile entry generation**

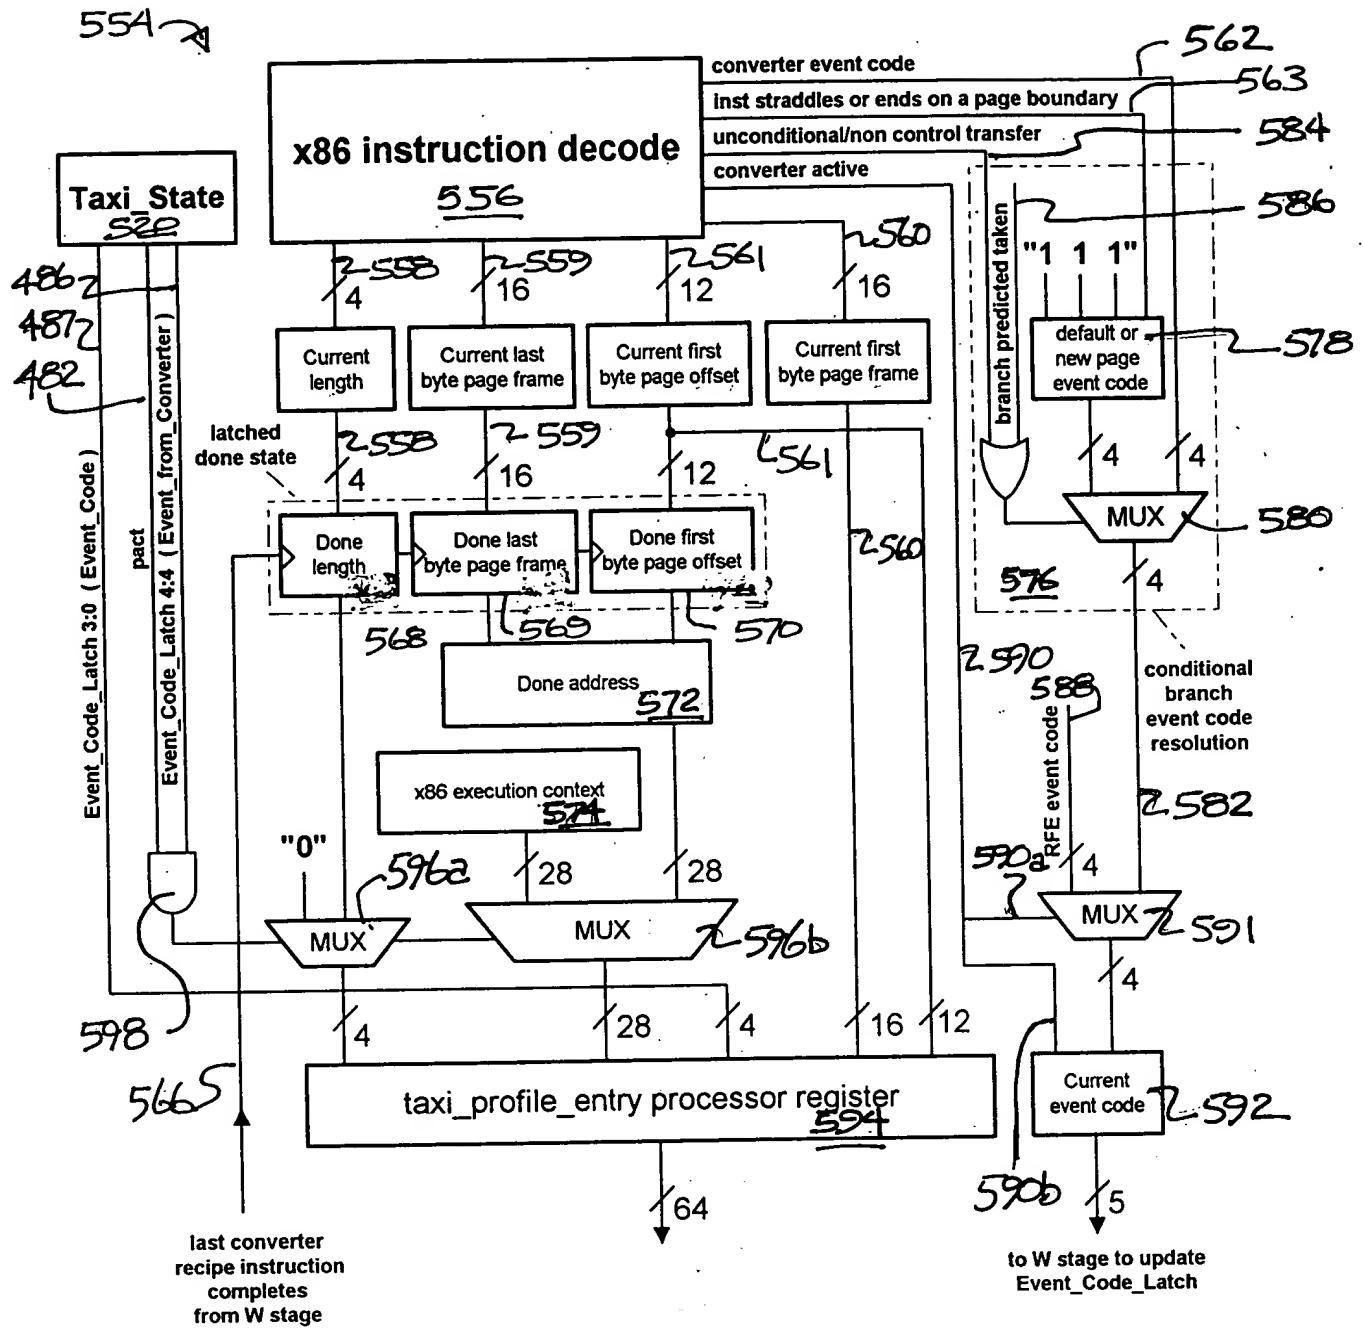


Fig. 5b

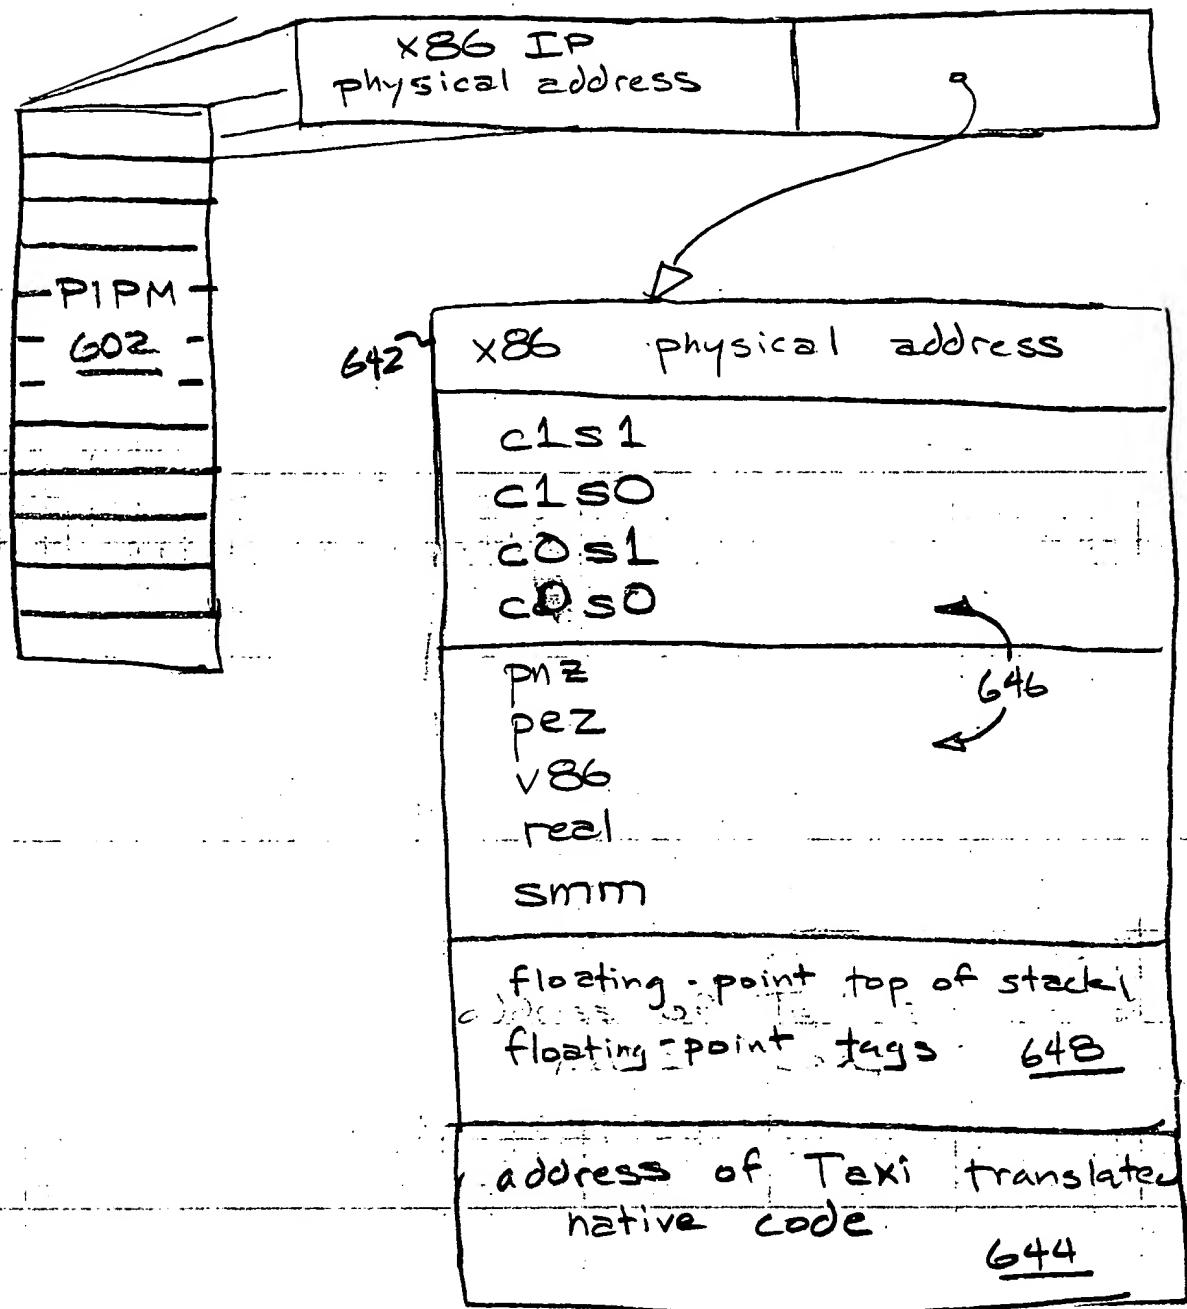


Fig. 6a

Event code from RFE restarting converter  
or mapping of converter's x86 opcode

RFE or previous converter cycle

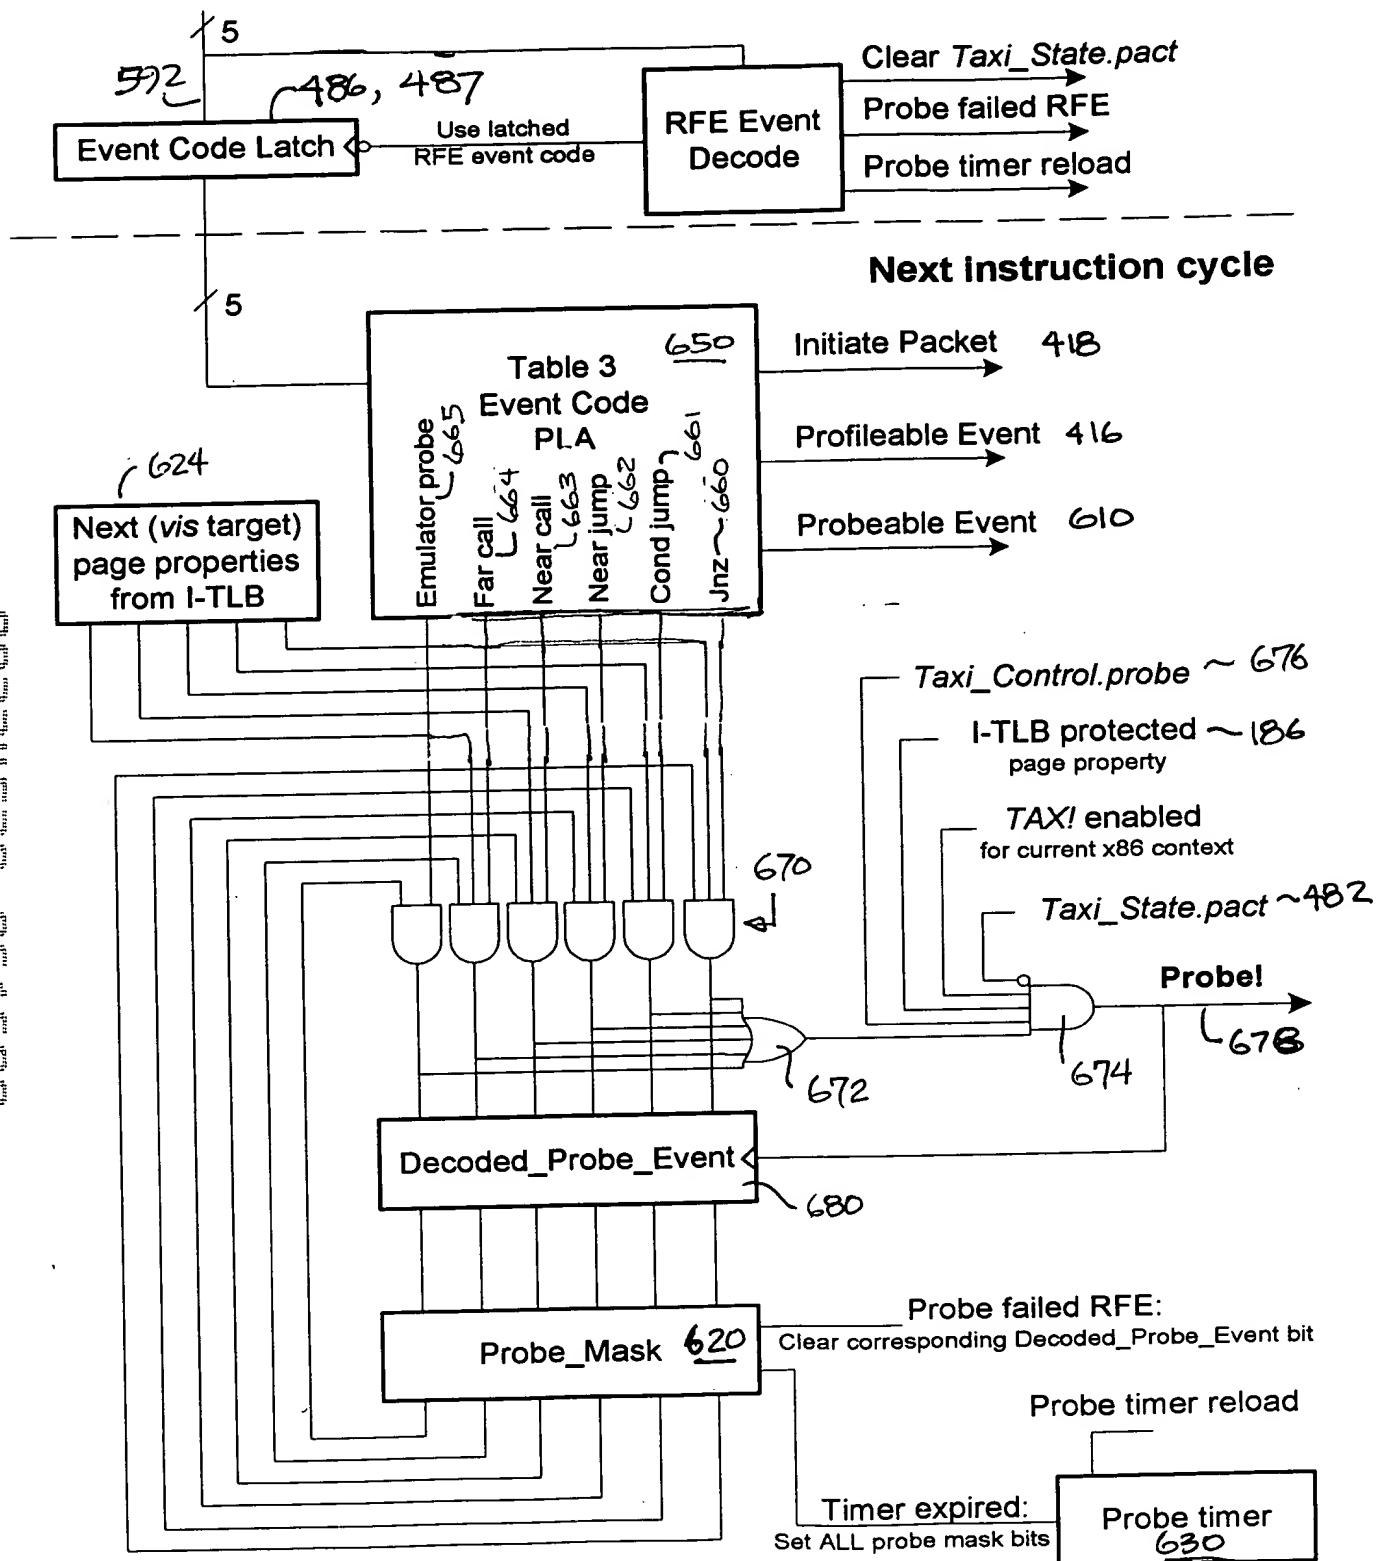


Fig. 6b

As each event occurs during execution of an X86 program in converter 136 or emulator 316, materialize an event code in event code latch 486, 487

650: PLA 650 processes the event code to produce at most one of five classifications of the event, "jnz" 660, "conditional jump" 661, "near jump" 662, "near call" 663, "far call" 664, or "emulator probe" 665

670: The bit 660-665 is ANDed with the probe page properties 624 from TLB 116 and Taxi\_State.Probe\_Mask 620

672: OR together the products of the ANDs. The sum of the OR represents the predicate "the event code 592 is an event on a page whose probeable event bit is currently enabled in Taxi\_State.Probe\_Mask 620 and the TLB copy of the PFAT page properties."

674: AND the sum of the OR together with several machine context predicates to see if this is a probeable event

690: Consult the bit vector to verify that the probeable event is in an address range with a corresponding translated code segment

682: Execute a TAXI instruction to materialize a Context\_At\_Point entry describing the current machine state, to supply arguments to the probe exception handler

Deliver a probe exception to transfer control to the software exception handler

Probe PIPM 602 for an entry 640 corresponding to the address of the target of the event

was a PIPM entry found?

N

mismatch

Evaluate/verify the preconditions from integer portion 686 of PIPM 602 entry 640

match

Evaluate/verify the preconditions from floating-point portion 688 of PIPM 602 entry 640, and if mismatching, unload floating-point context and reload it to conform to PIPM

Transfer control to the TAXI translated native code

Fail: resume execution  
of x86 binary in  
converter 136

Fig. 6c